

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-06-0008

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, creating and maintaining the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing burden, to AFRL, Attention: Office of Management and Budget, Paperwork Reduction Project and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project and Reports, Washington, DC 20503.

| | | | |
|--|---|---|----------------------------|
| 1. AGENCY USE ONLY (Leave blank) | 2. REPORT DATE | 3. REPORT TYPE AND Final | |
| 4. TITLE AND SUBTITLE Protocolol, Engineering Research Center | | 5. FUNDING NUMBERS f49620-00-1-0330 | |
| 6. AUTHOR(S) Dr. Melliar-Smith | | | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of CA, Santa Barbara | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR-Dr. Sjogren | | 10. SPONSORING/MONITORING AGENCY REPORT NUMBER | |
| 11. SUPPLEMENTARY NOTES | | | |
| 12a. DISTRIBUTION AVAILABILITY STATEMENT <i>Public release is approved.</i> | | 12b. DISTRIBUTION CODE | |
| 13. ABSTRACT (Maximum 200 words) During the period of this grant, Jennifer Hou moved from Ohio State University to the University of Illinois at Urbana Champaign. Professor Hou's portion of the MUM grant was transferred from Ohio State University to the University of Illinois at Urbana Champaign. She was promoted from Associate Professor to Professor in August 2005. Mary Baker left Stanford University to join Hewlett Packard Laboratories in 2003. Elizabeth BeldingRoyer at the University of California, Santa Barbara, replaced Mary Baker on this grant. She was promoted from Assistant Professor to Associate Professor in July 2005. Douglas Schmidt moved from the University of California, Irvine, to become a program manager at DARPA. Subsequently, he moved to Vanderbilt University where he is now a professor. Raymond Klefstad replaced Douglas Schmidt at the University of California, Irvine, as a principal investigator on this grant. | | | |
| 14. SUBJECT TERMS | | 15. NUMBER OF PAGES | |
| | | 16. PRICE CODE | |
| 17. SECURITY CLASSIFICATION OF REPORT | 18. SECURITY CLASSIFICATION OF THIS PAGE | 19. SECURITY CLASSIFICATION OF ABSTRACT | 20. LIMITATION OF ABSTRACT |

MURI/AFOSR FINAL REPORT

Protocol Engineering Research Center, University of California, Santa Barbara
Grant Number F49620-00-1-0330, June 2000 - December 2005

1. PI Information

P. M. Melliar-Smith Professor and Principal Investigator
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
Phone: 805-448-8250, Fax: 805-893-3262
Email: pmms@ece.ucsb.edu

Mary G. Baker Assistant Professor
Department of Computer Science
Stanford University
Phone: 650-725-3711, Fax: 650-725-2588
Email: mgbaker@cs.stanford.edu

Elizabeth Belding-Royer Associate Professor
Department of Computer Science
University of California, Santa Barbara, CA 93106
Phone: 805-893-3411, Fax: 805-893-8553
Email: ebelding@cs.ucsb.edu

J. J. Garcia-Luna-Aceves Professor
Department of Computer Engineering
University of California, Santa Cruz, CA 95064
Phone: 831-459-4153, Fax: 831-459-4829
E-mail: jj@cse.ucsc.edu

Jennifer C. Hou Professor
Department of Computer Science
University of Illinois at Urbana Champaign, Urbana 61801
Phone: 217-265-6329, Fax: 217-244-6500
E-mail: jhou@cs.uiuc.edu

Raymond Klefstad Assistant Adjunct Professor
Department of Electrical Engineering and Computer Science
University of California, Irvine, CA 92697
Phone: 949-824-1901, Fax: 949-824-2321
Email: klefstad@uci.edu

Karl N. Levitt Professor
Department of Computer Science
University of California, Davis
Davis, CA 95616
Phone: 530-752-2149, Fax: 530-752-4767
E-mail: levitt@cs.ucdavis.edu

Louise E. Moser Professor
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106
Phone: 805-893-4897, Fax: 805-893-3262
E-mail: moser@ece.ucsb.edu

Douglas C. Schmidt Professor
Department of Electrical and Computer Engineering
University of California, Irvine CA, 92697
Phone: 314-935-4215, Fax: 314-935-7302
Email: schmidt@cs.wustl.edu

Tatsuya Suda Professor
Department of Information and Computer Science
University of California, Irvine, CA 92697
Phone: 949-824-4105, Fax: 949-824-2886
Email: suda@ics.uci.edu

20060119 006

During the period of this grant, Jennifer Hou moved from Ohio State University to the University of Illinois at Urbana Champaign. Professor Hou's portion of the MURI grant was transferred from Ohio State University to the University of Illinois at Urbana Champaign. She was promoted from Associate Professor to Professor in August 2005.

Mary Baker left Stanford University to join Hewlett Packard Laboratories in 2003. Elizabeth Belding-Royer at the University of California, Santa Barbara, replaced Mary Baker on this grant. She was promoted from Assistant Professor to Associate Professor in July 2005.

Douglas Schmidt moved from the University of California, Irvine, to become a program manager at DARPA. Subsequently, he moved to Vanderbilt University where he is now a professor. Raymond Klefstad replaced Douglas Schmidt at the University of California, Irvine, as a principal investigator on this grant.

2. Project Objectives

The scientific goal of this AFOSR MURI effort is to advance the state-of-the-art in design, implementation, and evaluation of network protocols to meet the needs of complex distributed systems. More specifically, the research has focused on (i) the design and evaluation of protocols for real-time performance, fault tolerance and security, as well as protocols for mobile wireless environments, (ii) the design and implementation of analysis/simulation tools to validate/evaluate protocols, and (iii) the development of strategies and middleware to ease the programming of complex networked applications.

The ultimate objectives are to

- Design network protocols and environments that allow the programmers to create their applications without having to worry about timing, fault tolerance, mobility or security;
- Develop network protocols whose behavior is predictable, even under adverse conditions or in the absence of complete information;
- Develop design/analysis/verification tools that help developers of defense and other applications to design their network-centric systems and to understand effects of their design decisions; and
- Devise methodologies that allow application designers and programmers to express, to other designers, to toolsets, or to operating systems, the need of applications in terms of functionality, real-time performance, and fault-tolerant operations.

In this AFOSR MURI grant, the principal investigators focused their efforts on eight different research topics, namely:

- Protocols for Supporting Real Time Resource Management
- Protocols and Infrastructure for Supporting Fault Tolerance
- Methodology and Protocols for Supporting Security
- Protocols for Supporting Scalable Real-Time Multicasts
- Protocols for Supporting Mobile Wireless Environments
- Bio-Networking Architectures
- Middleware for Supporting Application Programs with End-to-End QoS
- Tools for Supporting Network Simulation/Emulation.

We summarize the major accomplishments, scientific research, and impact related to these research topics below. We also report on education and student training, publications, presentations, other professional activities, technology transfer, patent disclosures, software releases, and honors during the period supported by this AFOSR MURI grant.

3. Major Accomplishments

3.1 Protocols for Supporting Real Time Resource Management

- Enhanced the Real Time CORBA support in the TAO Object Request Broker (Schmidt).
- Designed and implemented an infrastructure, based on the TAO CORBA ORB, developed by Schmidt and Suda at UC Irvine, that integrates load balancing and fault tolerance (Melliar-Smith and Moser).
- Leveraged long-range dependency characteristics in Internet traffic, exploited its correlation structure for predicting future traffic at multiple time scales, utilized prediction results online for resource control in active queue management and congestion control in TCP (Hou).

3.2 Protocols and Infrastructure for Supporting Fault Tolerance

- Designed and implemented an interception-based fault-tolerant CORBA infrastructure, called the Eternal system. The Eternal system intercepts the IIOP messages that would have been sent by the CORBA ORB and instead multicasts them to the replicas of an object using a reliable totally-ordered multicast protocol (Melliar-Smith and Moser).
- Designed and implemented a pluggable fault-tolerant CORBA infrastructure, called the Pluggable Fault-Tolerant infrastructure, that provides fault tolerance for CORBA applications by utilizing the Pluggable Protocols Framework (Melliar-Smith and Moser).
- Developed the Totem Redundant Ring Protocol that supports three different network replication strategies: Active, Passive, and a hybrid Active-Passive network replication strategy (Melliar-Smith and Moser).
- Designed and implemented the Aroma system, which provides transparent fault tolerance for Java applications that use Java Remote Method Invocation (Moser and Melliar-Smith).
- Designed and implemented a software infrastructure that unifies transactions and replication for three-tier architectures and, thus, provides high availability and fault tolerance for enterprise applications. (Moser and Melliar-Smith).
- Designed and evaluated a reservation protocol strategy for business activities that span multiple enterprises, such as Web Services (Moser and Melliar-Smith).

3.3 Methodology and Protocols for Supporting Security

- Developed a formal framework to analyze and improve detection rules of intrusion detection systems. Validation of the model and the framework are established using the System Health and Intrusion Monitoring (SHIM) intrusion detection system (Levitt).
- Developed a formal framework to analyze and improve the specification-based intrusion detection rules developed for an auto-configuration protocol, i.e., the Dynamic Registration and Configuration Protocol (DRCP), used by mobile ad hoc networks (Levitt).
- Developed a distributed, automated response model that utilizes a Proportional-Integral-Derivative controller to aid in handling network traffic flow management. This response model is designed to address Distributed Denial of Service attacks that exploit the availability of servers and routers, resulting in the severe loss of their connectivity (Levitt).

- Developed a specification-based intrusion detection system for detecting attacks on the Ad Hoc On Demand Distance Vector (AODV) Routing protocol, designed for mobile ad hoc networks (Levitt).
- Developed a specification-based intrusion detection system for detecting attacks on the Optimized Link State Routing (OLSR) protocol, designed for mobile ad hoc networks (Levitt).
- Developed an intrusion monitor for the Dynamic Host Configuration Protocol (DHCP). This monitor is a plug-in preprocessor for Snort that detects actual behavior of network elements violating the security specifications that were developed in studying the behavior of DHCP (Levitt).
- Developed a general automatic response model for mobile ad hoc networks (Levitt).
- Developed a distributed event-driven message exchange model for intrusion detection within a mobile ad hoc networking environment (Levitt).
- Investigated the completeness and accuracy properties of unreliable fault detectors for asynchronous distributed systems that are subject to Byzantine faults (Moser and Melliar-Smith).
- Won the Wilkes Award from the British Computer Society for the paper, “Unreliable Byzantine Fault Detectors for Solving Consensus,” published in the Computer Journal (Moser and Melliar-Smith).
- Development of a classification model for attacks against ad hoc networks. Our model is represented as an attack cuboid. The axes of the cuboid are based on the security goals, the network layer the attack targets, and the type of attack. All currently known attacks can be classified using this model (Belding-Royer).
- Development of AODVSTAT, a tool for detecting local attacks in ad hoc networks. Currently, five different attacks can be detected by intrusion detection sensors running AODVSTAT (Belding-Royer).
- Development of Distributed AODVSTAT, a distributed tool for performing global intrusion detection in an ad hoc network. IDS agents running D-AODVSTAT communicate and correlate the information they collect to detect more general attacks against the network. Currently, the false propagation of messages attack can be detected using this tool (Belding-Royer).
- Testing and performance evaluation of AODVSTAT and D-AODVSTAT within an ad hoc network testbed. The testbed consists of eight laptops. The six defined attacks were implemented against this network in a series of experiments. Results show that over 98% of the attacks are detected with approximately 2% false positives (Belding-Royer)
- Development of the emulation system that enables the evaluation of AODVSTAT and D-AODVSTAT in large, diverse networks. Four attacks were simulated in networks as large as 50 nodes in a series of experiments. A variety of different movement patterns and network sizes were evaluated (Belding-Royer)
- Development of the Authenticated Routing for Ad hoc Networks (ARAN) protocol that detects and protects against malicious actions by third parties and peers in ad hoc environments (Belding-Royer).
- Design of a classification scheme that identifies different environments/scenarios and their associated security requirements (Belding-Royer).

3.4 Protocols for Supporting Scalable Real-Time Multicasts

- Produced designs of new scalable, real-time video multicast protocols, and produced simulation models and simulators, to investigate characteristics of the protocols (Suda).
- Designed a general temporal QoS framework for QoS-driven shared tree multicast routing and devised eligibility tests for member join/leave procedures (Hou).
- Developed a multicast group communication engine and bridge that allows the use of multiple group communication protocols concurrently (Moser and Melliar-Smith).
- Based on insights derived from experimental measurements of the Pluggable Fault-Tolerant Infrastructure, we also developed a new replication strategy using semi-active replication and low-latency real-time multicasts (Melliar-Smith and Moser).
- Won the best paper award at the International Symposium on Performance Evaluation of Computer and Telecommunication Systems for the paper, “End-to-end latency analysis and evaluation of a fault-tolerant CORBA infrastructure” (Melliar-Smith and Moser).

3.5 Protocols for Supporting Mobile Wireless Environments

- Investigated participation incentives for proper routing in ad hoc networks. In particular, we have developed two general types of routing enhancements to detect and handle misbehaving nodes (Baker).
- Examined the queuing dynamics of nodes in ad hoc networks across a wide mobility spectrum in both on-demand and pro-active protocols and studied how to improve the performance of ad hoc networks through innovative packet scheduling (Baker).
- Developed a number of novel schemes to ensure loop freedom in on demand routing protocols (Garcia-Luna-Aceves).
- Designed and analyzed a new type of routing protocols for ad hoc networks that combines on demand and proactive routing modalities (Garcia-Luna-Aceves).
- Developed a new analytical model for MAC protocols in multihop ad hoc networks applicable to omni-directional and directional antennas (Garcia-Luna-Aceves).
- Developed a new MAC protocol for ad hoc networks based on directional antennas that permits network nodes to schedule their transmissions to avoid collisions (Garcia-Luna-Aceves).
- Designed and analyzed the Hybrid Activation Multiple Access (HAMA), a new type of medium access control protocol for ad hoc networks capable of scheduling the transmission of unicast and broadcast traffic using two-hop neighborhood information (Garcia-Luna-Aceves).
- Developed a novel transmission scheduling approach for ad hoc networks based on two-hop neighborhood information that is fair and can be applied to node or link activation schemes (Garcia-Luna-Aceves).
- Developed a novel scheme for topology control in ad hoc networks that is based on two-hop neighborhood information (Garcia-Luna-Aceves).
- Developed novel schemes for broadcasting of information in ad hoc networks (Garcia-Luna-Aceves).
- Devised a Minimum Spanning Tree topology control algorithm, called Local Minimum Spanning Tree, and a variation (to support k-connectivity in heterogeneous environments) for wireless multi-hop networks, and analytically proved several important properties of this algorithm (Hou).
- Studied the issue of power management and investigated the tradeoffs between reducing the energy consumption and minimizing the performance penalty. Explored the design space and devised a

suite of power management protocols from asynchronous to synchronous, reactive to proactive, and agnostic to cross-layer information to utilizing hints from the higher layers (Hou).

- Performed a rigorous study of performance limits for wireless sensor networks with respect to coverage, connectivity, lifetime, and power, and designed protocols that aim to achieve these performance limits (Hou).
- Produced designs of new protocols for sensor networks, and produced simulation models and simulators, to investigate characteristics of the protocols (Suda).
- Developed a natural language interface and voice-enabled applications for mobile wireless personal handheld devices (Melliar-Smith and Moser).

3.6 Bio-Networking Architecture

- Implemented in software the Bio-Networking Architecture and its associated protocols and empirically evaluated the Bio-Networking Architecture protocols (Suda)
- Produced designs of protocols for discovery in peer-to-peer systems, based on the Bio-Networking Architecture, and produced simulation models and simulators, to investigate characteristics of the protocols (Suda).

3.7 Middleware for Supporting Application Programs with End-to-End QoS

- Developed a pluggable protocol framework for the TAO ORB (Schmidt).
- Added support for key QoS-enabled network protocols (including IntServ, DiffServ, and SCTP) to TAO's pluggable protocols framework (Schmidt).
- Designed and implemented Java-based CORBA middleware, called ZEN, that is highly optimized for real-time performance. ZEN has a modular and pluggable architecture that enables the configuration of transport protocols, object adapters, and ORB core, in particular for small code size (Klefstad and Schmidt).
- Developed RTZen, an RTSJ (Real-Time Specification for Java) implementation of real-time CORBA and the successor of ZEN. In addition to providing the benefits of ZEN, it provides much higher predictability and throughput. We also discovered design patterns for facilitating the development of RTSJ-based middleware and developed a tool that visualizes the scoped hierarchies of complex applications and that also locates memory leaks (Klefstad).
- Developed ZEN-kit, a graphical tool for customizing RTZen. This customization is achieved through modularizing the middleware so that features may be inserted or removed based on the application requirements (Klefstad).
- Evaluated the performance of Java-based CORBA Component Model technology for real-time and embedded systems (Klefstad).
- Received the Best Paper Award at the 36th IEEE Hawaii International Conference on System Science for our paper titled "Design and performance of a dynamically configurable, messaging protocols framework for real-time CORBA" (Klefstad and Schmidt).
- Designed and implemented an infrastructure, based on the TAO CORBA ORB, developed by Schmidt and Suda at UC Irvine, that integrates load balancing and fault tolerance (Melliar-Smith and Moser).

Produced designs of protocols for inter-domain resource exchange (iREX), and produced simulation models and simulators, to investigate characteristics of the protocols (Suda).

3.8 Tools for Supporting Network Simulation/Emulation

- Developed simulation techniques for very large application-level distributed protocols and began implementation of Narses using network models of varying levels of details (Baker).
- Completed the implementation, evaluation, and public release of a component-based, compositional network simulation and emulation environment, called J-Sim. J-Sim supports network simulation in both wired and wireless environments, including wireless sensor networks (Hou).

4. Scientific Research

To achieve the project objectives, we have worked synergistically on the following research topics:

4.1 Protocols for Supporting Real Time Resource Management

Exploitation of LRD in Internet Traffic for Resource and Traffic Control (Hou): In several studies, Internet traffic has been shown empirically to be long-range dependent and self-similar. These characteristics imply Internet traffic exhibits bursty behaviors across multiple time scales and, hence, network resources (such as bandwidth and buffer space) have to be *over-provisioned* according to peak traffic load (rather than average traffic load) in order to provide QoS. On the flip side, these characteristics imply that there exist correlation structures across multiple time scales that can be leveraged in traffic prediction and network planning. Central to the exploitation of the abundant correlation structure for traffic control and measurement are (i) the prediction of future traffic on multiple timescales based on recent measurements, and (ii) use of the prediction results in the decision making of traffic control/measurement. We have carried out research along two major directions. First, we investigated both fractional-model-based and non-fractional-model-based predictors and determine which predictor(s) render the best performance in predicting Internet traffic. The *fractional brownian motion* model, the *fractal ARIMA* model, and the *multiplicative wavelet* model are representatives of the former, and the *linear minimum mean square error* predictor fall in the latter. This study employed performance measures of accuracy, ease of deployment, computational complexity, and adaptability. It was based on analytical reasoning, ns-2 simulation, and empirical experiments with the large amount of traffic measurement and traces available from the Cooperative Association for Internet Data Analysis. Second, we have investigated how to harness traffic prediction (obtained with appropriate choices of traffic predictors) in traffic control/measurement for the three research thrust areas discussed below.

Leveraging Traffic Predictability in Active Queue Management (Hou): We have shown that the correlation structure present in long-range dependent traffic can be detected online and used to predict future traffic. We then devised a novel scheme, called *TCP with traffic prediction*, that exploits the prediction results to infer, in the context of AIMD steady-state dynamics, the optimal operational point at which a TCP connection should operate. Through analytical reasoning, we have shown that the impact of prediction errors on fairness is minimal. We have also conducted ns-2 simulation and FreeBSD 4.1 implementation studies to validate the design and to demonstrate the performance improvements in terms of the packet loss ratio and throughput attained by connections.

Exploration of the Use of Traffic Predictability to Improve TCP Congestion Control (Hou): With the objective of stabilizing the instantaneous queue length, we have incorporated the prediction results in the calculation of the packet dropping probability in active queue management (AQM). We call the resulting scheme Predictive AQM (PAQM). Through analytical reasoning, we have shown that PAQM is a generalized version of the well-known AQM scheme, random early detection (RED) that takes the future arrival rate as a new dimension of congestion index. By stabilizing the queue at a desirable level with consideration of future traffic, PAQM enables the link capacity to be fully utilized, while not incurring excessive packet loss ratio. Using ns-2 simulation, we have compared PAQM against existing AQM schemes with respect to different performance criteria, and have shown that under most cases PAQM outperforms RED in stabilizing the instantaneous queue length and adaptive virtual queue in reducing the packet loss ratio and achieving higher utilization of the link capacity.

Exploitation of LRD for Resource and Traffic Control (Hou): We have investigated three theoretically grounded packet-pair methods, prediction, reconstruction and interpolation, for measuring cross traffic on the bottleneck link of an end-to-end path. Essentially back-to-back, closely-spaced packet pairs are sent at a sender, and the inter-arrival times of packets in each pair are measured at the receiver and are used to infer the amount of cross traffic at the bottleneck link. The objective is to infer cross traffic as accurately as possible, while not injecting a significant number of probe packets into the network. To achieve this objective, we take advantage of the LRD characteristic of cross traffic and take measurements (i.e., send back-to-back packet pairs) only at time instants strategically determined by LMMS prediction, signal reconstruction, and interpolation. We have conducted simulation/empirical (Internet) studies to study (i) whether these methods give good mean or instantaneous measurements of cross traffic, and (ii) whether they are adaptive to dynamic changes of cross traffic and are robust in the presence of multiple bottleneck links on an end-to-end path.


Real Time CORBA Support (Schmidt): Developers of mission critical distributed real-time applications have historically used low-level, non-standard network programming interfaces, application-specific protocols, and customized real-time scheduling mechanisms. This legacy approach mixes real-time network programming throughout application programs, which makes their development non-portable, tedious, and error-prone, and requires highly skilled developers. To address these problems, we have developed and optimized the TAO Object Request Broker (ORB). TAO is an open-source, standards-compliant implementation of the Real Time CORBA specification that is being used on hundreds of research projects and commercial products world-wide. TAO is designed to meet end-to-end application QoS requirements by vertically integrating distributed object computing middleware with OS I/O subsystems, network protocols, and network interfaces.

TAO was the first real-time ORB end system to support end-to-end QoS guarantees over high-speed networks (such as ATM), embedded system interconnects (such as VME and Fibrechannel), and QoS-enabled network protocols (such as IntServ, DiffServ, and SCTP).

The following is a synopsis of the contributions of the Real Time CORBA portion of the PERC project:

- An ORB core that supports deterministic real-time scheduling and dispatching strategies. The TAO ORB core concurrency models minimize context switching, synchronization, dynamic memory allocation, and data movement.
- An active demultiplexing strategy that associates client requests with target objects in constant time, regardless of the number of objects and operations.
- A highly-optimized CORBA GIOP protocol engine and an IDL compiler that generates compiled or interpreted stubs and skeletons, that enables applications to make fine-grained time/space tradeoffs.


Load Balancing and Fault Tolerance (Melliar-Smith and Moser): Aditya Singh, Michael Melliar-Smith and Louise Moser developed an infrastructure, based on the TAO CORBA ORB, developed by Schmidt and Suda at UC Irvine, that integrates load balancing and fault tolerance.



Integration of fault tolerance and load balancing, where peer groups are used for load balancing and replica groups are used for fault tolerance.

4.2 Protocols and Infrastructure for Supporting Fault Tolerance

Interception-Based Fault Tolerant CORBA Infrastructure (Moser and Melliar-Smith): Professors Moser and Melliar-Smith and their students have designed and implemented the Eternal system that provides fault tolerance for CORBA applications using the interception approach. The Eternal system intercepts the IIOP messages that would have been sent by the CORBA ORB and, instead, multicasts them to the replicas of an object using a reliable totally-ordered multicast protocol. The Eternal system filters duplicate operations, and controls the dispatching of operations to the multiple threads of an object. It logs messages and checkpoints, and transfers the state of a replica to a new or recovering replica, and logs and replays messages.



The components of the Eternal system.

Design and Implementation of a Pluggable Fault Tolerant CORBA Infrastructure (Moser and Melliar-Smith): Wenbing Zhao, Louise Moser and Michael Melliar-Smith have developed a Pluggable Fault-Tolerant CORBA infrastructure (Pluggable FT) that provides fault tolerance for CORBA applications by utilizing the Pluggable Protocols Framework (PPF) that is available for most CORBA ORBs. The fault tolerance mechanisms, except for the Totem group communication protocol, are integrated into the application process's address space using the PPF. The approach does not require modification to the CORBA ORB, and requires only minimal modification to the application. Moreover, it avoids the difficulty of retrieving and assigning the ORB state, by incorporating the fault tolerance mechanisms into the ORB.

The Pluggable Protocols Framework separates the messaging and network protocols from other parts of the ORB core and from the application objects. The PPF allows both the network protocol and the messaging protocol to be replaced. Thus, the PPF makes it possible to develop CORBA applications for, and to deploy them in, environments for which the standard IIOP protocol is not appropriate. The PPF also provides CORBA applications with improved quality of service by enabling the use of customized protocols that are tailored to those applications and their environments.

In the Pluggable Protocols Infrastructure, we use the PPF to replace the network protocol with the Totem reliable totally-ordered multicast protocol. The infrastructure achieves performance that is better than that of existing Fault Tolerant CORBA systems, while providing strong replica consistency and supporting a wider range of applications.

Totem Redundant Ring Protocol (Melliar-Smith and Moser): Ruppert Koch, Michael Melliar-Smith and Louise Moser have developed the Totem Redundant Ring Protocol (Totem RRP) that supports three different network replication strategies: Active, Passive and a hybrid Active-Passive network replication strategy.

In active replication, all messages and tokens are sent over all networks at the same time. All data are received multiple times. The bandwidth consumption increases linearly with the number N of networks. The maximum throughput equals the throughput of the slowest network. The system is able to mask the loss of a message on up to N-1 networks without any message retransmission delay. In passive replication, messages are sent alternately over one of the available networks. The bandwidth consumption equals the bandwidth consumption of an unreplicated system. In the fault-free case, the maximum throughput equals the sum of the throughputs of all networks. If one of the networks fails, the maximum throughput is reduced. If a

message is lost, Totem must wait until the message has been retransmitted. Active-passive replication is a mixture of active replication and passive replication. Every message or token is sent over K networks simultaneously, $1 < K < N$. The bandwidth consumption increases K -fold. The system is able to mask the loss of a message on up to $K-1$ networks without any message retransmission delay.


Performance measurements show that a system using the Totem RRP with Active replication experiences a decrease in effective throughput, as is to be expected. In contrast, in a system that uses the Redundant Ring Protocol with Passive replication, the throughput exceeds the throughput of the unreplicated system while being more resilient to network faults.

Latency of Replication and Multicast Protocols (Melliar-Smith and Moser): Wenbing Zhao, Louise Moser and Michael Melliar-Smith have performed experimental measurements of the Pluggable Fault Tolerant CORBA Infrastructure and a reliable totally-ordered multicast group communication protocol. We measured probability density functions for the latency from a request transmitted by a client, through service by a (possibly replicated) server, to reply delivered to the client. These measurements show that, while the logical token ring protocol does achieve excellent throughput, it has an adverse effect on latency. The 120 microsecond latency measured for an unreplicated server is increased to 600 to 900 microsecond for a three-way replicated server. With an alternative rotating sequencer protocol, the latency is still about 400 microseconds for a three-way replicated server. This increased latency is not significant in enterprise applications operating over a wide-area network, but is undesirable in real-time embedded applications.

Based on insights derived from these measurements, we have developed a new replication strategy using semi-active replication. Messages are processed immediately at the primary replica and that message order is communicated to the backup replicas so that they can mimic the behavior of the primary replicas. Considerable care is required to ensure that a faulty processor coupled with the loss of messages cannot leave the system in an inconsistent state. Preliminary measurements show a latency of 170 microseconds, which we hope to reduce to about 150 microseconds.

Hub Implemented Protocols (Melliar-Smith and Moser): Ruppert Koch, Michael Melliar-Smith and Louise Moser have investigated novel implementation strategies for fault tolerant protocols. There have been recent efforts to move the protocol stack from the host computer to the network interface card. We are investigating moving parts of the protocol stack one step further to the Ethernet switch. It is still necessary to operate a simple protocol between the switch and the host computer, currently TCP, but many of the difficult issues are handled more easily at the switch, including multicasting, message ordering, retransmission, and flow control. The primary challenge is whether it is possible to achieve adequate performance from the Ethernet switch while it performs these additional functions. Another challenge is replication of the Ethernet switch to avoid a single point of failure, while still ensuring consistency of the state of the two switches.

Fault Tolerance for Three-Tier Architectures (Moser and Melliar-Smith): Wenbing Zhao, Louise Moser and Michael Melliar-Smith have developed a software infrastructure that unifies transactions and replication for three-tier architectures and, thus, provides high availability and fault tolerance for enterprise applications. The infrastructure is based on the Fault Tolerant CORBA and CORBA Object Transaction Service standards. Within the middle-tier at the server, the infrastructure replicates the applications to protect the business logic processing. In addition, it replicates the transaction coordinator, which renders the two-phase commit protocol non-blocking and, thus, avoids potentially long service disruptions caused by failure of the coordinator. The infrastructure implements automatic client-side failover mechanisms, which guarantee that clients know the outcome of the requests that they have made. Moreover, it handles the interactions between the applications and the database system through replicated gateways that prevent duplicate requests from reaching the database system. The infrastructure starts the transactions at the applications, and automatically retries aborted transactions, caused by process or communication faults, on behalf of the clients.



Evaluation of Fault Detection and Recovery Times (Melliar-Smith and Moser): Xiaoli Chen, Michael Melliar-Smith and Louise Moser developed a benchmark, with characteristics generally similar to NSWC's HiPerD air defense application. The benchmark consists of a driver and a measurement process on one computer, and a four-stage pipeline with each stage implemented on a different computer. Each of the four stages was replicated using Eternal System's real-time fault tolerance replication infrastructure.

Measurements were made for the time to fault detection and recovery with all operations performed at the same priority level, and with two priority levels thus allowing a small proportion of the operations to be high priority operations. In our experiments, it became apparent immediately that the primary determinant of the fault detection and recovery time was the 10 ms resolution of the Linux clock.

The time through the four-stage pipeline was mean 2.654 ms, maximum 12.867 ms, without replication and fault tolerance, and mean 3.783 ms, maximum 12.441 ms, with primary/backup replication and fault tolerance. For high priority operations, the average time was 3.878 ms, the maximum 4.909 ms. The fault detection and recovery time was 138 milliseconds, which is too long. With a faster clock, the fault detection and recovery time would have been shorter.

TCP Connection Failover for Fault Recovery (Melliar-Smith and Moser): Ruppert Koch, Sanjay Hortikar, Louise Moser and Michael Melliar-Smith developed a protocol that allows failover of a TCP server endpoint in a transparent manner. The failover can occur at any time during the lifetime of the connection, and is achieved by modifying the server's TCP/IP stack. No modifications are required to the client application, the server application, the client's TCP/IP stack, or the network.

The essence of the problem is that failover must be achieved at one end of the TCP connection without disturbing the other end, which might be external to the system and thus might not be running any TCP failover software. In particular, the new TCP server endpoint, taking over after the failure of the original TCP server endpoint, must continue using the same sequence numbers, acknowledge messages already received, and retransmit messages for which acknowledgments are not received.

Thus, the standby server TCP endpoint must receive every message that the primary server TCP endpoint transmits or receives, which is achieved by having the primary server TCP endpoint forward all such messages to the standby server TCP endpoint. Fortunately, the TCP stack contains a mechanism for inserting additional modules into the stack that can be used to achieve this effect. The TCP stack also contains mechanisms for setting, during startup, the initial sequence numbers, which can be used to ensure that the standby server TCP endpoint uses the correct sequence numbers when it takes over.

TCP connection failover is, of course, only an adjunct to failover of an application program. If the application program is using passive replication with restart from a checkpoint, the application program needs to have messages replayed from a log, and will generate messages that have already been sent and that must be suppressed. The TCP connection failover mechanisms at the standby server TCP endpoint include mechanisms for buffering and logging messages. These mechanisms must be integrated with message logging and replay mechanisms for the application.

Consistent File Replication (Melliar-Smith and Moser) : Rachit Chawla, Michael Melliar-Smith and Louise Moser investigated strategies for consistent file replication that maintain consistency between the states of files and the states of the application programs that use those files, thus simplifying recovery from faults. We have considered four application availability policies:

- Simple restart
- Checkpoint restore
- Passive replication
- Active replication.

For each of these policies, we have investigated corresponding file replication strategies.

Consistent Time Service for Fault-Tolerant Distributed Systems (Melliar-Smith and Moser): Clock-related operations are one of the many sources of replica non-determinism and of replica inconsistency in fault-tolerant distributed systems. In passive replication, if the primary server crashes, the next clock value returned by the new primary server might have actually rolled back in time, which can lead to undesirable consequences for the replicated application. The same problem can happen for active replication when the result of the first replica to respond is taken as the next clock value.


Wenbing Zhao, Louise Moser and Michael Melliar-Smith have designed and implemented a Consistent Time Service for fault-tolerant distributed systems. The Consistent Time Service introduces a group clock that is consistent across the replicas and ensures the determinism of the replicas with respect to clock-related operations. The group clock is monotonically increasing, transparent to the application, and fault-tolerant. The Consistent Time Service guarantees the consistency of the group clock even when faults occur, new replicas are added to the group, and failed replicas recover.

A Reservation Protocol for Business Activities That Span Multiple Enterprises (Moser and Melliar-Smith): Traditional transaction semantics are not appropriate for business activities that involve long-running transactions in a loosely coupled distributed environment, in particular, for Web Services that operate between different enterprises over the Internet. Extended transaction models have been developed to coordinate business activities, but they rely on compensating transactions to achieve data consistency and atomicity of tasks in the presence of faults. Even though the state change resulting from a particular committed task might be reversed by a compensating transaction, it is difficult to cancel the side effects of a committed task because other transactions might see the results before the compensating transaction is applied. Identifying such transactions and compensating all of them is difficult.

Wenbing Zhao, Louise Moser and Michael Melliar-Smith have developed a novel reservation-based extended transaction protocol that can be used to coordinate such business activities. The protocol avoids the use of compensating transactions, which can result in undesirable effects. In our protocol, each task within a business activity is executed as two steps. The first step involves an explicit reservation of resources. The second step involves the confirmation or cancellation of the reservation. Each step is executed as a separate traditional short-running transaction. We have investigated how our protocol can be implemented as a reservation protocol on top of the Web Services Transaction specification or, alternatively, as a coordination protocol on top of the Web Services Coordination specification.

Analysis of the Reservation Protocol Strategy (Melliar-Smith and Moser): We have conducted several availability and performance analyses in order to evaluate the reservation protocol strategy in comparison with the abort of transactions with compensating transactions strategy and the distributed transactions with two-phase commit strategy.

The probability that the databases are left in a potentially inconsistent state is an important availability metric. The figure shows the probabilities of potential inconsistency for (a) the abort of transactions with compensating transactions strategy, and (b) the reservation protocol strategy. The reservation protocol has superior performance because there are fewer additional transactions for each fault recovery. It is our assessment that the difference in the probabilities that the database is left in a potentially inconsistent state presents a decisive advantage for the reservation protocol strategy. If the distributed transactions with two-phase commit strategy is used, instead of the abort of transactions with compensating transactions strategy, the risk of inconsistency is reduced but there is an increased risk that a fault will cause data to be locked for an arbitrarily long period of time.




The probability that the database is left in a potentially inconsistent state after m business activities for (a) the abort of transactions with compensating transactions on the left, and (b) the reservation protocol strategy on the right.

The probability density functions (pdfs) for the duration of a business activity with delays due to lock contention is also an important metric. The figure shows the pdfs for the duration of a business activity for different values of lock contention for both transactional locking and the reservation protocol. When the probability of lock contention is low, the pdfs for the duration of a business activity using transactional locking are substantially as expected, and the effects of delays due to contention for a single lock and for two locks are clearly visible. As the probability of contention for a lock increases, the business activities are delayed, locks are held longer, delays due to lock contention are longer, and the probability that a business activity claims a lock that is already held by another business activity increases. The resulting pdfs have long tails and, thus, there is a high probability of lengthy delays for the business activity. It is worth noting that, for transactional locking, there are probabilities for lock contention for which the system is not stable, representing unbounded delays and essentially no progress for the business activity. Such lock contention and instability lead to system collapse under heavy load, which often occurs at the most inappropriate times during the most important tasks. This observation underscores the importance of determining, and enforcing, an admission control limit for the business activities.

Also shown in the figure are pdfs for the duration of a business activity for the reservation protocol. It is evident that even high probabilities of contention for locks do not result in substantial delays for the business activity, because locks are held only briefly during the reservation subtask and are not held for the full duration of the business activity. In summary, the reservation protocol strategy is more resilient to high loads and high probabilities of lock contention than the transactional locking strategy.

If, instead of distributed transactions with two-phase commit, only local transactions and compensating transactions are used, the reservation of part of the resource is simpler. The concurrency achieved is substantially equivalent to that of the reservation protocol strategy.



The probability density functions for the duration of a business activity with delays due to lock contention.

The Heisenbug Problem (Melliar-Smith and Moser): Transient errors of unknown cause, sometimes called Heisenbugs, are a perplexing problem in computer systems. Heisenbugs are significantly more frequent, up to 30 times more frequent, than permanent faults or Bohrbugs, and are the primary determinant of the reliability of a computer system. Because Heisenbugs are difficult to reproduce, it is difficult to identify the cause of these errors and to remove them from the computer system.

Michael Melliar-Smith and Louise Moser have investigated four strategies for addressing the Heisenbug problem:

- Acceptance/checkpoint strategy
- Transactional strategy
- Strategy employing consistent diversity in time
- Strategy employing consistent diversity in space

The most appropriate strategy to employ depends on the nature of the application and the environment in which it runs.

4.3 Methodology and Protocols for Supporting Security


Intrusion Detection for Mobile Ad Hoc Networks (Levitt). Our research in this area comprises a specification-based IDS for the AODV protocol, a specification-based IDS for the OLSR protocol, an intrusion detection architecture for MANETs, and a distributed event-driven message exchange model for MANET IDS. We discuss each of these topics in more detail below.

- **Specification-based IDS for the AODV Protocol.** The Ad-Hoc On-Demand Distance Vector (AODV) Routing protocol, designed for mobile ad hoc networks, offers quick adaptation to dynamic link conditions, low processing and memory overhead, and low network utilization. However, because the protocol design did not address security issues, AODV is vulnerable to various kinds of attacks. In our research, we analyzed some of the vulnerabilities, specifically addressing attacks against AODV that manipulate the routing messages, and have proposed a solution based on the specification-based intrusion detection technique to detect attacks on AODV. Briefly, our approach involves the use of finite state machines for specifying correct AODV routing behavior, and distributed network monitors for detecting run-time violations of the specifications. We propose one additional field in the protocol message to enable the monitoring. Our algorithm can effectively detect most of the serious attacks in real time and with minimum overhead.
- **Specification-based IDS for the OLSR Protocol.** The Optimal Link State Routing (OLSR) protocol, developed at INRIA, France, is another popular routing protocol. Unlike AODV, this protocol is a proactive table-driven link-state protocol that has been modified for the MANET environment. Local Hello beacons and a multi-point relay system are employed to obtain local state information at a node and disseminate it optimally around the network. In our research, we conducted an extensive analysis of the vulnerabilities in the OLSR routing protocol, and develop an OLSR IDS model. We established a set of security requirements that must be satisfied in protocol behavior. We established a set of detection rules that strictly enforce the identified security requirements during protocol execution. We use GloMoSim, a scalable simulation environment for ad hoc networks, to implement the detection rules and example attacks, and to identify false alarms and exceptions. This simulation is coupled with ongoing formal reasoning and proof methodology to establish that the specifications satisfy the security policy under various assumptions.
- **Intrusion Detection Architecture for MANETs.** Intrusion detection in MANETs is challenging because these networks change their topologies dynamically, lack concentration points where aggregated traffic can be analyzed, utilize infrastructure protocols that are susceptible to manipulation, and rely on noisy, intermittent wireless communication. We have developed a cooperative, distributed intrusion detection architecture that addresses these challenges while facilitating accurate detection of MANET-specific and conventional attacks. The architecture is organized as a dynamic hierarchy in which detection data is acquired at the leaves and is incrementally aggregated, reduced, and analyzed as it flows upward toward the root. Security management directives flow downward from the nodes at the top. To maintain communication efficiency, the hierarchy is automatically reconfigured as needed using clustering techniques in which clusterheads are selected based on topology and other criteria. The utility of the architecture is illustrated via multiple attack scenarios.
- **Distributed Event-driven Message Exchange Model for MANET IDS.** Numerous distributed intrusion detection systems have been proposed to detect abnormal routing behaviors in MANET. Most of these systems simply assume that each detector has complete information from peers instead

of building an efficient message exchange framework among detectors. To solve this problem, we propose a Distributed Event-driven Message Exchange Model (DEMEM) which provides sufficient data for intrusion detection in MANET. In DEMEM, detectors residing in each mobile node intercept routing messages of the node and exchange information with other detectors only when they have new routing information. This event-driven approach can avoid both costly promiscuous monitoring and modifying routing protocols. We demonstrated DEMEM on a specification-based intrusion detection model that precisely detects insider routing attacks in Optimized Link State Routing, while only requiring 2-hop neighbors to exchange their Hello messages. DEMEM has additional ID Request messages to eliminate false positives caused by lost messages. We validated our approach with formal analyses showing that DEMEM ensures sufficient information is provided for each detector in a distributed OLSR IDS. In addition, we conducted simulation experiments in GloMoSim demonstrating low message overhead, no false negatives and very low false positives.

Formal Reasoning and Verification (Levitt). In this research area, we have performed formal reasoning and verification about intrusion detection systems, and also about a specification-based intrusion detection for dynamic auto-configuration protocols in ad hoc networks.

- **A Specification-Based Intrusion Detection for Dynamic Auto-Configuration Protocols in Ad Hoc Networks.** In this research topic, we have addressed security issues of auto-configuration protocols in ad hoc networks. Auto-configuration protocols enable nodes to obtain configuration information so that they can communicate with other nodes in the network. We have described a formal approach to modeling and reasoning about auto-configuration protocols to support the detection of malicious insider nodes. Our approach defines a global security requirement for the subnet with respect to the protocol and involves the analysis of the local detection rules that characterize the good behavior of individual nodes. The local detection rules define a distributed specification-based intrusion detection system to detect malicious insider nodes. Novel to the field of intrusion detection is a formal proof that the local detection rules ensure the global security requirement. Our proof, in principle, can detect any attack, even unknown attacks, that can imperil the global security requirement. Our proof rests on assumptions that reflect security-imperiling behavior that they rule out; the assumptions, which can be strengthened or relaxed as the threat changes, simplify the intrusion detection system to attacks that are unlikely.



Hierarchical framework for verification.

- **Intrusion Detection Systems.** We have developed a formal framework for the analysis of intrusion detection systems (IDS) that employ declarative rules for attack recognition, e.g., specification-based intrusion detection. Our approach allows reasoning about the effectiveness of an IDS. A formal framework is built with the theorem prover ACL2 to analyze and improve detection rules of IDSs. System Health and Intrusion Monitoring (SHIM) is used as an example specification-based IDS to validate our approach. We formalized specifications of a host-based IDS in SHIM which, together with a trusted file policy, enables us to reason about the soundness and completeness of the specifications by proving that the specifications satisfy the policy under various assumptions. These assumptions are properties of the system that are not checked by the IDS. Analysis of these assumptions shows the beneficial role of SHIM in improving the security of the system. The formal framework (see the figure) and analysis methodology provide a scientific basis for arguing that an IDS can detect known and unknown attacks because the IDS detects all attacks that violate a policy.

Automated Response (Levitt). Our research in this area comprises a network-based response framework and an automatic response model for MANETs. These topics are summarized below.

- **Network-Based Response Framework.** Distributed Denial of Service (DDoS) attacks exploit the availability of servers and routers, resulting in the severe loss of connectivity. We have developed a distributed, automated response model that utilizes a Proportional-Integral-Derivative controller to aid in handling traffic flow management. This model is designed to prevent incoming traffic from exceeding a given threshold, while allowing as much incoming, legitimate traffic as possible. In addition, this model focuses on requiring less demanding modifications to external routers and networks than other published distributed response models that impact the effect of DDoS attacks.
- **Automatic Response Model for MANETs.** We have addressed the problem of automatic response to attacks on MANETs, which is a challenging problem because the topologies of these networks are dynamic and the fully distributed cooperation among nodes makes not only the false positives of intrusion alarms high but also the response analysis complicated. We have developed a cooperative, distributed and automatic response model that addresses these challenges by correlating intrusion alarms and then deciding which response among a set of feasible responses is the best one to take. Our work focuses on automatic response to “atomic” routing message attacks of insiders, based on an attack taxonomy that we developed to provide structure to the ways an insider can deleteriously impact routing. Four basic response types can be effected in our system: isolation, relocation, encryption and routing recovery. The decision on which response to take is based on the criticality of the node’s position in the topology and on the attack. To determine the importance of a node’s position in a network, we have proposed a new concept, topology dependency, to represent the positional relationship between each pair of neighboring nodes. The utility of this general response model is illustrated via numerous atomic attack scenarios.

Worm Defense (Levitt). There is a great need for a way to reproduce faithfully live environments for worm- and worm-defense research. In our research, we have made use of a network testbed called EMULAB, to satisfy this need and evaluate our framework for worm-defense evaluation within enterprise networks, i.e., networks with a few tens to a few thousands of nodes. We have described an implementation of the framework and use it to evaluate an example defense strategy, but emphasize that the framework can support many different defense strategies. The framework is encapsulated in an API. This API accepts a topology description and a description of the defense system, and evaluates the defense system against worms. The worms can be characterized by a specification or operationally by a worm program.

Unreliable Byzantine Fault Detectors for Solving Consensus (Moser and Melliar-Smith). Unreliable fault detectors can be defined in terms of completeness and accuracy properties, and can be used to solve the consensus problem in asynchronous distributed systems that are subject to crash faults.

Kim Kihlstrom, Louise Moser and Michael Melliar-Smith have extended this result to asynchronous distributed systems that are subject to Byzantine faults. First, we defined and categorized Byzantine faults. We then defined two new completeness properties, eventual strong completeness and eventual weak completeness. We used these completeness properties and previously defined accuracy properties to define four new classes of unreliable Byzantine fault detectors. Then, we developed an algorithm that uses a Byzantine fault detector to solve the consensus problem in an asynchronous distributed system of n processes in which the number k of Byzantine faults satisfies $k \leq (n-1)/3$. We also developed algorithms that implement a Byzantine fault detector in a model of partial synchrony. We proved the correctness of the consensus algorithm and analyzed its complexity.


Checkpointing and Logging for Intrusion Analysis and Recovery (Moser and Melliar-Smith). Intrusions often occur because of operating system and application program vulnerabilities. It is important to analyze thoroughly how an intrusion occurred so that the related vulnerabilities can be found and fixed quickly. Unfortunately, audit records produced through traditional system logging do not contain enough information to provide certainty on how the system was compromised and what damage was done. Such logs are even less useful to recover from an intrusion.

Recently, researchers have started to employ techniques that have been used in the fault tolerance community, such as logging of external inputs and non-deterministic events in order to replay the execution of the software, to facilitate better intrusion analysis. Wenbing Zhao, Louise Moser and Michael Melliar-Smith have designed a software infrastructure for intrusion analysis and recovery. In this infrastructure,

intrusion analysis is facilitated by deterministic logging and replay. Recovery from intrusion is carried out by rolling back to the last checkpoint before the intrusion occurred.

The AODVSTAT Intrusion Detection Framework (Belding-Royer). Because a complete security system needs both intrusion prevention and intrusion detection techniques, we have developed an intrusion detection framework, called AODVSTAT, that provides intrusion detection to ad hoc networks running the AODV routing protocol. AODVSTAT can detect intrusions both locally and globally within the ad hoc network. Experiments have shown that our solution provides effective intrusion detection functionality using a limited amount of resources. Our documented version of AODVSTAT is now freely available for download on the Internet.

As another part of this work, we developed a classification model for attacks against ad hoc networks. We use a three-dimensional model in the form of a cuboid to classify attacks in an ad hoc network. Across the x-axis, we group attacks based on the type of threat they pose to the network. Along the y-axis, each attack is classified based on the layer at which the attack can be launched. Along the z-axis, attacks are classified according to the behavior of the nodes and the type of the attack. The primary behaviors demonstrated by attackers are silence, selfishness, misbehavior, mischief and interference. The classification scheme for attacks that we have developed is shown in the figure below.



The classification scheme for attacks that we developed.


The AODVSTAT Emulation System (Belding-Royer). To expand the capabilities of AODVSTAT and to test it in larger systems, we have developed an emulation system for AODVSTAT integrated with the NS-2 network simulator. IDS sensors are placed within the simulated network. When they receive messages, these messages are captured and then recreated as actual Ethernet frames so that they can be sent to the AODVSTAT module for processing. The emulation system allows us to use the actual AODVSTAT code in the evaluation of large, diverse networks. The emulation system enables use to investigate the performance of large networks, the size of which is only limited by the processing power of the machines running the simulations. To date, we have studied networks as large as 50 nodes.

The Authenticated Routing for Ad hoc Networks (ARAN) Security Protocol (Belding-Royer). The ARAN protocol detects and protects against malicious actions by third parties and peers in managed-open ad hoc environments. ARAN introduces authentication, message integrity and non-repudiation to an ad hoc environment as a part of a minimal security policy. Our evaluations show ARAN has minimal

- Performance costs for the increased security in terms of processing
- Networking overhead.


As another part of this work, we defined a set of three discrete ad hoc wireless environments based the opportunity for pre-deployment exchange of security parameters. These environments are termed open, managed-open and managed-hostile. By clearly defining the level of security needed in a particular scenario, the appropriate security mechanisms can be deployed to protect the network while not needlessly burdening it with security procedures.

The figure below shows a graph of the average packet delivery fraction of ARAN and AODV. The graph indicates that, even though ARAN has more overhead due to the security mechanisms, it is still able to successfully deliver the same percentage of packets as AODV.



The average packet delivery fraction of ARAN and AODV.

The next figure shows a graph of the fraction of data packets sent through malicious nodes with ARAN and AODV. It indicates that, when there are malicious nodes in the network, those nodes are able to influence the routes AODV selects such that a larger percentage of packets traverse those nodes. Malicious nodes can snoop, modify and drop packets, as well as perform any of a number of other undesirable events. However, with ARAN, malicious nodes have no way of influencing the paths that are selected, and so only if they are on the shortest path between the source and destination will they end up on the selected path. Hence, far fewer packets end up traversing malicious nodes.



The fraction of data packets sent through malicious nodes with ARAN and AODV.

4.4 Protocols for Supporting Scalable Real-Time Multicasts

QoS-Driven Multicast Routing (Hou): We have investigated the problem of QoS provisioning in shared multicast routing protocols such as Core-Based Tree, Simple Multicast and Protocol Independent Multicast Sparse Mode, and devised a unified QoS extension framework for these kinds of multicast routing protocols. Specifically, the QoS requirements are heterogeneous among receivers, and are specified as either (R1) an upper/lower bound of an additive, multiplicative, or concave QoS parameter (e.g., the end-to-end delay bound), or (R2) an upper bound of the inter-receiver difference of a QoS parameter (e.g., the inter-receiver delay jitter bound). Considering these QoS requirements, we have devised a set of enhancements to the member join/leave and state update/refresh procedures to facilitate the deployment of additive (e.g., end-to-end delay bound), multiplicative (e.g., packet loss ratio along a path) and concave (e.g., minimum bandwidth available) QoS. In particular, we have (i) devised eligibility tests to check whether a new member can join a multicast tree at adequate QoS, without violating the existing QoS to the other members on the tree, (ii) determined the set of states kept at each router in order to conduct eligibility tests, (iii) devised a state update/refresh procedure that is based on soft state and can be readily integrated with the tree maintenance mechanism that already exists in most core-based multicast routing protocols, and (iv) implemented, based on the CBT daemon in FreeBSD (developed by British Telecom), and empirically evaluated the performance of the QoS framework. We have reported our evaluation results in the Inter-Domain Multicast Routing Working Group of the Internet Engineering Task Force.

Scalable, real-time video multicast protocols (Suda): In this research, we have investigated a novel protocol, called Source-Adaptive Multi-layered Multicast (SAMM), to control congestion caused by real-time video multicast over the Internet. SAMM relies on the exchange of feedback between end systems (a video source and multicast destinations) and rate adjustment at a video source. In SAMM, multi-layered video encoding is deployed, and raw video data is encoded into one or more layers of differing priority. Video sources adapt the number of video layers they generate as well as the transmission rate of each video layer in response to congestion feedback from the network and receivers. Using simulations that incorporate multi-layered video codecs, we have demonstrated that SAMM exhibits better scalability and responsiveness to congestion than existing protocols. Accomplishments in this work include design of the SAMM protocols, a simulator to evaluate the proposed protocols and simulation results confirming the performance of the proposed protocols.

Multicast Group Communication Engine and Bridge (Moser and Melliar-Smith): Multicast group communication protocols are needed for fault-tolerant distributed systems to maintain strong replica consistency. However, different multicast protocols are appropriate for different applications and environments. As part of this MURI project, we have defined a multicast group communication engine and bridge that allows use of multiple group communication protocols concurrently. The group communication engine uses Lamport timestamps for message ordering and heartbeat (null) messages for liveness. Timestamps and heartbeat messages provide a convenient mechanism by which multiple group communication protocols can coexist without constraining the on-the-wire message formats and internal algorithms of the protocols

The group communication engine places timestamps on messages, and multicasts messages to groups, using one or more group communication protocols. Each multicast group communication protocol reliably delivers its timestamped messages in timestamp order to the group communication engine. The group communication engine integrates the streams of messages into a single stream of messages for delivery in timestamp order. No forwarding or conversion from one multicast protocol to another is necessary, and the multicast protocols never communicate directly with one another or even know of each others' existence.

Real-Time Fault Tolerance Protocol (Melliar-Smith and Moser): Wenbing Zhao, Michael Melliar-Smith and Louise Moser have developed a Real-Time Fault Tolerance Protocol that provides fault tolerance for distributed applications using the primary/backup process replication approach. The protocol suite consists of a reliable messaging protocol, a membership protocol, and a virtual determinizer. The messaging protocol provides a reliable ordered message delivery service by employing a direct group-to-group multicast, where the ordering is determined by the primary replica in each process group. The membership protocol provides a fast reconfiguration and recovery service when one or more replicas become faulty and when a replica joins or leaves a process group. The virtual determinizer captures the ordering information at the primary and enforces the same ordering at the backups for several major sources of replica non-determinism.

4.5 Protocols for Supporting Mobile Wireless Environments

Mitigating Misbehavior in Ad Hoc Routing Protocols (Baker): Gaining reasonable traffic throughput in ad hoc networks requires the nodes in the network to participate in packet forwarding on behalf of other nodes. The presence of nodes that misbehave in this respect degrades network performance. Misbehaving nodes can assume at least two forms: nodes that agree to forward packets on behalf of others but then fail to do so, or merely selfish nodes that attempt to preserve their own resources by not agreeing to forward packets.

We have developed two general types of routing enhancements to handle such misbehaviors. The first set of techniques, called *Watchdog* and *Pathrater*, allows us to isolate and route around nodes that agree to participate but fail to do so. The second technique, called *OCEAN (Observation-based Cooperation Enforcement in Ad hoc Networks)*, attempts to limit both forms of misbehavior through punishing bad nodes by denying their traffic.

Our approach in both avoiding misbehaving nodes and in providing participation incentives is to avoid modifying existing routing protocols, which would be error-prone and complex. Instead, we provide information to an existing routing protocol to allow it to make a better choice than the default route choice.

A second goal, especially with the incentive-based techniques, is that the protocol enhancements be lightweight. This means that in forming opinions about the behavior of other nodes, we rely only on directly observed events, and not on reputation information as reported by other parties. Such direct observation allows us to avoid the overhead of propagating reputation information throughout the network. In addition, this means that unlike Blazevic's work we do not need a tamper-resistant layer on each node, and unlike Buchegger's work we do not need signing or other encryption.

We simulated Watchdog and Pathrater in the Dynamic Source Routing protocol (DSR) implementation in the ns-2 simulator. The *Watchdog* and *Pathrater* algorithms when used together in an ad hoc network of moderate mobility running the Dynamic Source Routing protocol (DSR) improve packet throughput by 17% in the presence of even 40% misbehaving nodes while increasing routing protocol overhead by 9% (from DSR's 12% to 17%).

Our incentive-based techniques are still under development. So far we have evaluated these techniques within a high-level Java simulation environment specifically designed for testing such protocols quickly. This simulator does not give us accurate timing results, since it has no model for the underlying wireless network, but it allows us to measure overall throughput during the course of the simulation. Within these simulations, we find that OCEAN improves the traffic flows of well-behaved nodes by up to 15% while dramatically lowering the throughput of misbehaving nodes (by up to 75%).

Improving Ad Hoc Network Performance through New Packet Scheduling Algorithms (Baker): We have investigated improving the overall performance of ad hoc networks through new packet scheduling. For this study, we examined the queuing dynamics of nodes in an ad hoc network across a wide mobility spectrum (from static to non-stop motion) in both on-demand (DSR) and pro-active protocols (Greedy Perimeter Stateless Routing, or GPSR [Karp]) as implemented in the ns-2 simulator. For this packet scheduling work, we again do not modify the routing protocols themselves but instead modify only the processing of packet queues at a node. We find that the common scheduling practice of giving priority to routing control packets over data packets is an advantage in on-demand routing protocols such as DSR, but that it can actually reduce performance in proactive routing protocols. Most ad hoc network packet schedulers do not distinguish between different types of data packets, but we find that it is useful to do so. In particular, giving priority to data packets with short distance metrics (fewest remaining hops in DSR and shortest remaining distance in GPSR) shows the smallest delay – about a 30% reduction over standard DSR and standard GPSR and the highest throughput without increasing routing overhead.

Wireless Networks (Garcia-Luna-Aceves): Garcia-Luna-Aceves and his students have designed new algorithms and protocols for channel access, routing, topology control and broadcasting; verified the correctness of these algorithms and protocols; and analyzed their performance through analytical models and simulations. The simulations were carried out in GloMoSim, Qualnet and ns2, which are used widely in the research community.

We have developed analytical models for collision avoidance protocols operating in multihop ad hoc networks with omni-directional or directional. These models were extensions of prior work on CSMA and ALOHA and the first to address collision avoidance schemes in multihop wireless networks.

We have developed new channel access protocols based on node and link activation and applied them to sensor networks. We have started the development of flow-oriented scheduled channel access, which extends our prior results on scheduled channel access by using the flows traversing network nodes the entities that compete for scheduled channel access.

In contrast to a wired network, the scheduling decisions made at the MAC layer in wireless networks impact the de-facto topology over which routing and multicasting must operate. Hence, it is important to understand the ability to manage the useful topology of a wireless network to make network-level and end-to-end protocols more effective in multihop wireless networks. We have developed a novel approach to the solution of the connected dominating set election problem, which we call Topology Management by Priority Ordering (TMPO). Our approach uses the neighbor-aware contention resolution (NCR) algorithm employed in our prior collision-free channel access protocols to provide fast convergence and load-balancing with regard to the battery life and mobility of mobile nodes. Based on NCR, TMPO assigns randomized priorities to mobile stations, and elects a minimal dominating set (MDS) and the connected dominating set (CDS) of an ad hoc network according to these priorities. In doing so, TMPO requires only two-hop neighbor information for the MDS elections. The dynamic priorities assigned to nodes are derived from the node identifiers and their willingness to participate in the backbone formations. The willingness of a node is a function of the mobility and battery life of the node. The integrated consideration of mobility, battery life and deterministic node priorities makes TMPO the best-performing heuristic for topology management in ad hoc networks to date.

We have developed a mesh-oriented multicast protocol that eliminates the limitations of tree-based multicasting, ODMRP, and CAMP. Our protocol, which we called PUMA is the first multicast routing protocol for ad hoc networks based on meshes and a receiver-initiated group joining scheme that does not require an underlying unicast routing protocol to operate or the pre-assignment of cores to groups. In a nutshell, the first receiver joining a group becomes the core of the group and starts transmitting core announcements periodically. Each such packet specifies a sequence number, the address of the group, the address of the core, the sending node, and the distance to the core. Routers use the best core announcements they receive to send their own core announcements to their neighbors, and over time each router has one or multiple paths to the core of each known group in the ad hoc network. To join a multicast group, a router sends a join announcement to its next-hop toward the core of the group, which it learns from core announcements. Nodes receiving join announcements intended for them join the group and also send a join announcement periodically to their next-hops for the group core. To attain a mesh structure, routers that overhear join announcement packets not addressed to them from two or more neighbors join the group silently, i.e., they do not send join announcements. Normal group members send join announcements periodically. A multicast data packet for a group is forwarded from its source towards the core of the group, using next-hop information obtained in core announcements, and is flooded within the mesh of the group as soon as it reaches the first mesh member. Members and silent members participate in the flooding of data packets within a mesh. We have shown through extensive simulations in QualNet that PUMA dramatically outperforms ODMRP and MAODV.


We have developed several techniques to maintain loop-free on-demand routing in ad hoc networks, regardless of the changes in the underlying network topology. These schemes address the use of path information as labels used for ordering of nodes with respect to a destination, the use of sequence numbers as ordering labels, the use of path information together with information about links that caused routes to change, and the combination of sequence numbers with limited inter-neighbor coordination. All the proposed schemes work on the basis of route requests (RREQ), route replies (RREP), and route errors (RERR) in their messaging structure, which is similar to that used in DSR and AODV. However, the proposed schemes use new loop-free invariants for each destination. All of these schemes have been shown to outperform AODV and DSR significantly. Furthermore, we have demonstrated that AODV and similar protocols based on sequence numbers can incur looping and counting to infinity problems when nodes delete routing state for certain destinations due to rebooting or other causes. AODV and the more recent DYMO specification can incorporate our proposed schemes to solve the looping problems we uncovered.

We have investigated distributed algorithms for pruning the number of nodes needed to forward route requests or other type of broadcast packets in ad hoc networks. We have also developed a number of algorithms to support clustering in ad hoc networks, such that a given node can be covered by multiple cluster heads, each within a maximum number of hops, while trying to minimize the number of cluster heads needed to cover all network nodes.


Topology Control in Mobile Ad-hoc Networks (Hou): Topology control and management (how to determine the transmit power of each node so as to maintain network connectivity while consuming the minimum possible power) is an important issue in mobile ad hoc networks (MANETs). Instead of transmitting using the maximum possible power, nodes in a MANET collaboratively determine their transmit power and define the topology of the MANET by their neighbor relations under certain criteria.

Topology control and management is important in that it critically affects the system performance in a cross-layer manner. Determination of the transmit power level affects several node attributes: (i) the transmission range (as well as the interference caused by transmission), and (ii) the node degree (and hence the scope of the physical neighborhood). Consequently, topology control and management influences the Physical layer through the quality of the received signal, the MAC layer through the interference and contention caused by communication, the Network layer through the set of links that are formed (and hence the connectivity), the Transport layer through the overall data transport capability (and hence the network capacity) and, last but not least, the drain on battery energy. As a result, determining adequate transmit power in a *decentralized* manner is critical to several aspects of sensor network functionality. It affects network spatial reuse and hence network capacity. It affects network connectivity -- choosing too large a transmit power level results in excessive interference, while choosing too small a transmit power level can result in a disconnected network. It also has an effect on the contention for the medium. MAC-level contention and collision can be mitigated as much as possible by choosing the smallest transmit power subject to maintaining network connectivity. Finally, it affects the energy use for communication, and thus impacts battery life, a critical resource for many MANET applications.

We have devised a Local Minimum Spanning Tree (LMST) algorithm and variations for topology control and management. In LMST, each node builds its local minimum spanning tree *independently* with the use of locally collected information, and keeps only ontree nodes that are one hop away from its neighbors in the final topology. We have proved analytically that, if every node executes LMST, the network connectivity is preserved (i.e., if the original topology, in which each node uses the maximal transmit power is connected, the topology derived under LMST is also connected). An important feature of LMST is that it is fully *localized* (i.e., it depends only on the information collected locally), and hence is less susceptible to the impact of mobility and incurs less message overhead/delay. In addition, we have proved that (1) the node degree of any node in the resulting topology is bounded by 6, and (2) the topology can be transformed into one with bi-directional links (without impairing network connectivity) after removal of all unidirectional links (see the first figure below). Simulation results indicate that, compared with existing topology control algorithms (e.g., cone-based topology control by Li et al. and relay-region-based topology control by Meng et al.), the topology derived under LMST has smaller average node degrees (both logical and physical) and smaller average link length (see the second figure below). The former reduces the MAC-level contention, while the latter implies that small transmission power is needed to maintain connectivity. Moreover, as shown in the third figure below, LMST realizes a system throughput improvement of 25-50%.




Network topologies with the use of maximum transmission.



Performance comparisons w.r.t. node degree and average length of links among different algorithms. n nodes are uniformly distributed in a 1 km x 1 km region, where n varies from 50 to 250.

As part of the extension to LMST, we have shown that most of the topology control algorithms (including LMST) render sub-optimal performance or even fail *when* different nodes have different maximal transmission ranges. The major problem is that the connection between a pair of nodes might not be bi-directional in a topology of heterogeneous networks. This scenario is typical in the battlefield where different warfare entities have different capabilities. To deal with this problem, we have extended LMST and devised a localized algorithm for heterogeneous wireless networks with non-uniform transmission ranges. Essentially, a minimum spamming graph (instead of a minimum spanning tree) is built independently in the neighborhood of each node.


Another extension of LMST that we investigated is to equip topology-controlled networks with fault tolerance capabilities. In a controlled topology, if a node fails (due to power depletion and/or malicious destruction) or moves away, the network is more susceptible to temporary disconnection (as a result of reduced routing redundancy). Obviously, there exists a tradeoff between route redundancy and other performance aspects (power consumption, spatial reuse, MAC level interference, and network capacity). We have extended LMST to preserve k -connectivity and have evaluated it via simulation and empirical studies with respect to the decrease in overall system throughput caused by incorporating fault tolerance.



Total throughput and energy efficiency.

The total amount of data delivered (in bytes) and the energy efficiency (in bytes/J). The n nodes are randomly distributed in a 1500m x 200m region, $n/2$ of which are sources and the other $n/2$ are destinations, with CBR traffic between them. The two-ray ground model is used as the propagation model, IEEE 802.11 is used as the MAC protocol, and AODV is used as the routing protocol. The start time of a connection is uniformly distributed in the range [25s, 50s] in a simulation run of 200 seconds.

Power Management for Mobile Ad-Hoc Environments (Hou): We have investigated the issue of power management in multi-hop wireless networks. The study hinges on the observations that wireless devices incur significant power consumption in idle states as demonstrated by other experimental studies and our own measurement results. As most modern wireless devices can be set to different power states to conserve energy, power management (which turns devices to low-power states when they are not in use) is a promising technique to conserve energy. The objective of this study was to devise and evaluate power management schemes to reduce the energy consumption in idle periods while minimizing the performance penalty incurred.



Performance trade-off between end-to-end delay and energy efficiency.

In this study, we discovered that the effectiveness of power management is determined by the degree of coordination among communication entities and the amount of cross-layer information utilized. We have

devised a suite of protocols that explore the design space from asynchronous to synchronous, reactive to proactive, and from agnostic to cross-layer information to utilization of hints from the higher layers. We have quantified the impact of the degree of coordination and pro-activeness on energy-performance tradeoffs through combinatorial analysis and queuing models. Finally, we evaluated the protocols using both event-driven simulations and experiments on a lab motes testbed. We have designed the communication stack and modular support for power management in TinyOS and also a suite of benchmark tests to profile the energy consumption of wireless sensor nodes. Specifically, we have carried out the following research tasks:

- Reactiveness in maintaining and distributing state information is an important design principle in multi-hop wireless communication. On the one hand, distribution of state information can be costly in terms of power consumption due to the size of the network. On the other hand, state information can become stale quickly because of network dynamics. We have applied such design principles and have proposed an on-demand power management framework to address the second problem. The key idea is to switch a node to low-power state upon detection of prolonged idleness. Soft states for power management are established and refreshed by data, and control packet transmissions. On timeout of the soft state, a node is switched to the low-power state. We have conducted extensive simulations to study the sensitivity of the parameter settings, the impact of cross-layer information and the performance of on-demand power management in conjunction with various routing protocols. We have demonstrated that the on-demand power management can indeed effectively tie the energy consumption with the communication in the network, as only nodes on the communication path are kept active (see the figure below). Although the reactive protocols usually suffer from long startup time in establishing the initial set of states, on-demand power management alleviates such problems by providing a control parameter to adjust the delay incurred with respect to traffic load.
- Power management in multi-hop wireless networks differs from stand-alone systems in that distributed coordination is required, as all the entities have to be put in the active state for the communication to take place. An important question is how to coordinate the power management states among distributed communication entities without global synchronization. Without prior knowledge of the set of neighbors and in highly mobile environments where neighbor information is volatile, it is desirable to have deterministic bounds to discover the neighbors and their power management schedules while minimizing the energy consumption in idle periods. We have designed an asynchronous wakeup schedule based on the theory of block designs. The idea is similar to ensuring mutual exclusion in distributed systems but with a stronger shift-invariant property. The theoretical result is applicable to both homogeneous and heterogeneous scenarios where nodes have similar or different residual battery power. The asynchronous wakeup mechanism works with on-demand power management seamlessly to provide energy savings with a controllable impact on performance. From the simulation studies, we observe 40-80% savings in energy consumption, while keeping the average delay comparable to a network without power management.
- It has long been recognized that energy conservation usually comes at the cost of degraded performance such as delay and throughput both in stand-alone systems and communication networks. We have developed analytical models to quantify such a tradeoff under different power management policies. Based on the decision when to put nodes to low-power states, we have categorized power management policies into two classes, time-out driven and polling-based. A M/G/1/K model with multiple vacations and attention spans is used to characterize the timeout driven policies, while a transient analysis is applied to derive state transition probabilities in polling-based systems. We have shown numerically how different system parameters affect the energy-performance tradeoff. We found that for time-out driven power management policies, the optimal policy exhibits a threshold structure, i.e., when the traffic load is below a certain threshold, a node should switch to the low-power state whenever possible and remain active otherwise. From our analysis, contrary to popular beliefs, polling-based policies such as the IEEE 802.11 PSM are not energy efficient for light traffic load.
- To evaluate empirically the performance of our proposed power management schemes, we have devised power management modules and a communication stack in TinyOS with emphasis on the generality and robustness of the communication primitives supported. The set of APIs can be used

for various routing and power management protocols. Then we implemented the on-demand power management framework and the asynchronous wakeup mechanism on a motes testbed. We have also carried out a suite of benchmark tests to profile the energy consumption of sensor nodes in TinyOS.

Fundamental Performance Limits of Wireless Sensor Networks (Hou): Driven by advances in MEMS micro-sensors, wireless networking, and embedded processing, ad-hoc networks of devices and sensors with (limited) sensing and wireless communication capabilities are becoming increasingly available for commercial and military applications such as environmental monitoring (e.g., traffic, habitat, security), industrial sensing, and diagnostics (e.g., factory, appliances), critical infrastructure protection (e.g., power grids, water distribution, waste disposal), and situational awareness for battlefield applications. Much has been written about how, once deployed, these wireless networks will affect the way we monitor environments, track objects, fight wars, and recover from disasters.

Interests in wireless sensor networks have indeed opened up new research areas and have led to a fairly large number of research activities in the areas of protocol design and system building/prototyping. However, *comparatively* little work has been done on understanding the fundamental performance limits of wireless sensor networks, e.g., the asymptotic behaviors of these networks with respect to network capacity, longevity, scalability, coverage, connectivity, and critical power required to maintain connectivity. In this study, we aimed to perform a rigorous study of performance limits for wireless sensor networks with respect to coverage, connectivity, lifetime, and critical power analysis. We have also designed the protocols that aim to approach these performance limits. Specifically, we have carried out several research tasks along the following lines:

- **Critical Power, Node Degree and Other Node Attributes for Maintaining Connectivity.** We have investigated the critical power required (in the almost surely sense) to maintain connectivity and how it scales as the network size or density increases. Once the power levels used by nodes are determined, their transmission range, neighbor relation, and other graph properties (such as clique number and chromatic number) are also determined. Thus, we have studied the relationship between these node/network attributes and connectivity. In particular, we have derived the critical node degree needed to maintain network connectivity, where the critical node degree is the smallest integer d such that the d -nearest-neighbor graph over the sensor network maintains connectivity. Also, we have investigated the extent to which topology control (i.e., the mechanism that allows each node to transmit using different power levels) conserves power (in the asymptotic sense), and have devised, based on these theoretical findings, localized topology control algorithms.
- **Optimal Conditions for k -coverage and Their Use in Devising Localized Algorithms.** After the sensing threshold (and hence the sensing range) is set, the area in which a sensor can perform its sensing task is determined. To prolong the network lifetime, it is desirable that a subset of nodes is selected (on a rotational basis) to provide k -coverage for at least \square -portion of the monitored area. We have explored optimality conditions based on which a subset of working nodes can be chosen to maintain k -coverage. We have devised, based on the optimality conditions, fully decentralized and localized algorithms for k -coverage in large-scale sensor networks.
- **Network Capacity and Lifetime.** Once the power level is set, and the coverage/connectivity condition is determined, the network capacity and lifetime can be analyzed. We seek the relationship between the minimum energy required and the network capacity that results. We believe the key is to model the locations of nodes, as well as the propagation attenuation, and then to study the transport capacity. We are also looking into the relationship between k -coverage and network lifetime and between k -coverage and node density. Given the lifetime T of a single wireless node, the network lifetime is upper bounded by T times the maximum value of k . With the above derived relationships, we can answer several fundamental questions such as (a) how many wireless nodes (the nodal density) have to be deployed in a region, in order to continuously monitor the region for a period of time kT , and (b) given the number of wireless nodes, what is the maximal possible lifetime that can be achieved by any algorithm?
- **Large-scale J-Sim Network Simulation.** J-Sim supports network emulation at the MAC level, which allows us to interface a set of real-life Berkeley mica motes with a large-scale sensor network (of desired size), seamlessly transport data between them, and carry out a high-fidelity scalability study. Currently, we are studying (a) whether the performance limits still hold when some of the assumptions (e.g., the Poisson point assumption, the Toroidal model, and the sensing/transmission

disk assumption) are not valid, and (b) whether, and beyond what network size, the asymptotic attributes reasonably characterize network behaviors.

Protocols for Sensor Networks (Suda). Sensor networks have emerged rapidly to provide surveillance functions in a variety of applications. With sensor networks, it is possible to provide uninterrupted surveillance of a large area without human intervention. Our sensor network research has focused on two topics, network coverage and data dissemination. On the topic of network coverage, we have investigated protocols that efficiently maintain surveillance of either the areas of interest or the monitoring targets of interest. We have three ongoing projects in this area, Protocol for Coverage-Aware Sensor Engagement in Sensor Networks, Protocol for Monitoring a Target from Multiple Viewing Angles in Sensor Networks, and Protocol for Protecting Targets in Sensor Networks. On the topic of data dissemination, we have investigated protocols that efficiently store sensing data in a network and disseminate the sensing data to users. We have one project in this area, Protocol for Data Dissemination in Sensor Networks. More details are provided below.

Protocol for Coverage-Aware Sensor Engagement in Sensor Networks (Suda). We have investigated a new protocol, Coverage-Aware Sensor Engagement (CASE), to efficiently provide required network surveillance of the areas of interest. The CASE protocol achieves the goal by selectively activating sensors according to the requirements of applications. To select which sensors should be activated, each sensor evaluates its own contribution to network surveillance in a distributed manner. Based on its contribution to network surveillance, each sensor determines if it should stay active and continue monitoring or if it should deactivate itself. Sensors with higher contribution to network surveillance preferably remain active, resulting in the need for fewer active sensors to guarantee enough surveillance.

We have conducted extensive simulations to evaluate the performance of the proposed CASE protocol in terms of the number of activated sensors and the communication and computation costs. The performance metrics are measured by varying sensor network density and the surveillance requirements of applications. Simulation results show that CASE, compared with existing work, guarantees the required network surveillance by a smaller number of active sensors while causing lower communication and computation costs. In addition, we have investigated the performance of CASE under the scenarios with location error and hotspot areas that need to be monitored with better surveillance. CASE presents enough robustness against location error and exhibits the ability to support hotspot areas.

Protocol for Monitoring a Target from Multiple Viewing Angles in Sensor Networks (Suda). We have investigated a protocol that selectively activates a group of sensors to form a monitoring structure around a mobile target such that activated sensors monitor the target from multiple viewing angles. The proposed protocol assumes that a large number of sensors with varying capabilities (e.g., different sensing and communication capabilities, and different mobility capabilities) are randomly deployed, creating a large-scale and dynamic sensor network environment. It also assumes that sensors can take one of the three different modes (i.e., sleeping mode, listening mode, and active mode). On detecting a mobile target, active sensors broadcast notifications to their neighboring sensors and activate neighboring sensors in listening mode in order to form a monitoring structure. Based on the information contained in the notifications (e.g., viewing angles to the target, distance to the target etc.), each activated sensor autonomously decides whether to stay in active mode or switch to sleeping mode. This distributed and autonomous activation process results in the self-organization of a monitoring structure that cannot be achieved by any individual sensor.

We have considered different scenarios based on different sensor capabilities and have developed a set of distributed protocols to selectively activate sensors and form an effective monitoring structure. We have designed simulation models to examine the set of protocols, and have developed an emulator to investigate the performance of the protocol running on real hardware. Simulation and emulation results confirm the effectiveness of the protocol.

Protocol for Protecting Targets in Sensor Networks (Suda). We have investigated a protocol to monitor and protect a mobile target in a large-scale and dynamic sensor network consisting of a large number of mobile sensors with varying capabilities. The proposed protocol enables mobile sensors to self-organize, without a central controller, forming a largest possible perimeter that surrounds a mobile target (or mobile targets) such that no intruder can penetrate the perimeter without being detected. In the design of the protocol, mobile sensors and mobile targets are modeled as particles that exist in the nature (such as molecules). The concept of repelling and attracting forces that exist among the particles in nature (such as inter-molecular forces) is introduced to determine the direction of the movement of sensors to form a

perimeter. By applying the inter-molecular force on mobile sensors, mobile sensors form a perimeter autonomously in a distributed manner.

We have designed a protocol that controls the movement of mobile sensors to form an optimal perimeter around mobile targets. We have conducted extensive simulations to examine the scalability and robustness of the protocol. Simulation results have verified the effectiveness of the protocol.


Protocol for Data Dissemination in Sensor Networks (Suda) We have investigated a new protocol, Rendezvous Track protocol, for sensors to determine where to store sensing data and for users to determine where to retrieve the sensing data. In the proposed protocol, each sensor applies a hash function to the attributes of the sensing data, yielding the location of a group of sensors to store the sensing data. When retrieving sensing data, a user applies the same hash function to the attributes of the desired sensing data to obtain the location of a group of sensors that store the desired sensing data. This group of sensors serves as rendezvous sensors for sensing data with some particular attributes so that sensing data and user queries containing the same attributes will always converge at the same set of rendezvous sensors, leading to successful sensing data retrieval. The P.I. also developed a mechanism through which the proposed protocol dynamically adjusts the group of sensors to store sensing data based on different user query patterns.

We have designed the Rendezvous Track protocol with dynamic adaptation to user query patterns. We have also designed and implemented simulation models to investigate the scalability, robustness, adaptability, and efficiency of the protocol. Simulation results have verified the effectiveness of the proposed protocol.

Spoken Natural Language Interface for a Personal Handheld Device (Melliar-Smith and Moser). Mobile wireless devices, such as cell phones, can provide for their users, potentially, many useful services but they are currently limited by inadequate input and output capabilities, a small screen and no or unusable keyboard and mouse. Karthik Perumalsamy, Rama Alebouyeh, Michael Schuricht, Yi Xin Hu, Shreyas Prasad, Zachary Davis, Michael Melliar-Smith and Louise Moser have experimented with speech input and output to augment the small screen and facilitate the provision of services to the user.

For speech synthesis, we used the NaturalVoices package developed by AT&T Laboratories, with excellent results. Speech output is a poor mechanism for communicating complex information, such as maps, lists, etc. Consequently, our applications use the screen to display such complex information. However, speech output is an excellent mechanism for responding to and confirming speech input.

For speech recognition, we used the DynaSpeak software developed at SRI International for DARPA. In the current state of the art, it is necessary to use a restricted vocabulary and grammar that is appropriate to a specific service or application. It is also necessary to train the recognizer for a specific user, the owner of the mobile device. This training can achieve much better recognition than is possible for call centers that must handle a wide range of users.



The Program Manager determines which service to invoke

We developed several natural language speech input applications for the mobile wireless handheld device, including a Contact Service, a Location Service, and an Appointment and Reminder Service. These services operate quite well, and switching between the services is reasonably natural and convenient.

The mobile wireless device must support multiple applications concurrently, and some single sentences involve multiple applications. For example, in the sentence

“Is there an Italian restaurant near Susan’s house?”


the phrase “... Susan’s house” is handled by the Contact Service, while the phrase “Is there an Italian restaurant near ...” is handled by the Location Service. Thus, the Program Manager initiates the process of speech recognition with a scan for keywords. DynaSpeak then reparses the sentence using a vocabulary and grammar appropriate to the service corresponding to the detected keywords. If keywords are detected for more than one service, the sentence is parsed separately for each of those services. The grammars contain wild cards that can match anything and, in particular, phrases that a different service should handle. While not instantaneous, multiple parses of a single sentence can be performed sufficiently rapidly to achieve reasonably natural and timely response to human speech input.

To develop and demonstrate this technology, we used a small handheld PC, manufactured by OQO. This device contains a 1 GHz Transmeta processor and 512 Mbytes of RAM, and is fast enough to run the speech recognition software. The device operates under the Windows operating system to provide a convenient development environment. The device has a poor microphone with no noise cancellation, which makes the device very sensitive to background noise. Moreover, the device has no loudspeaker and, thus, we employ either a Bluetooth earpiece or a WiFi link to a laptop computer with a loudspeaker. The WiFi link also provides the Internet access that several of the applications need.

4.6 Bio-Networking Architecture

The Bio-Networking Architecture, developed by Tatsuya Suda, is motivated by the observation that desirable properties of future networks (such as scalability, adaptability, survivability and availability) have already been realized in various biological systems. The Bio-Networking Architecture is a framework for developing large-scale, highly distributed, heterogeneous and dynamic network applications.

In the Bio-Networking Architecture, key biological principles and mechanisms are applied to design network applications. A network application is implemented as a decentralized collection of autonomous objects called *cyber-entities*, illustrated in the figure below. This structure is analogous to a bee colony (a network application) consisting of multiple bees (cyber-entities). Each cyber-entity implements a functional component related to its service or application, and follows simple behavior rules (e.g. replication, reproduction and migration) similar to biological entities.



Cyber-entities and network applications.

The key aspects of the Bio-Networking Architecture supported by this AFOSR MURI grant are:


- Investigation of Dynamic Service Composition
- Investigation of Adaptation and Evolution of cyber-entities
- Design and Implementation of the Bio-Networking Platform.

Major accomplishments in the investigation of the Bio-Networking Architecture are summarized below.

Dynamic Service Composition (Suda). Dynamic service composition, i.e., composing an application through autonomous interactions among multiple service components (i.e. cyber-entities) at runtime, is a key feature of the Bio-Networking Architecture. To enable dynamic service composition, we have proposed a new component model, Component Service Model with Semantics (CoSMoS). CoSMoS integrates the semantic information and the functional information of a service component into a single semantic graph representation. We have also developed a unified interface, Component Runtime Environment (CoRE), in

order to convert different implementations of service components into the CoSMoS representation. Using the semantic support of CoSMoS, we have developed a semantics-based service composition protocol, Semantic Graph based Service Composition (SeGSeC). SeGSeC generates the workflow of the application requested by a user and examines the semantics of the workflow against the user's request. See the figure below.

Using CoSMoS, CoRE and SeGSeC, we have implemented a service composition system and demonstrated that the implemented service composition system supports the semantics-based dynamic service composition. We have also evaluated the performance of SeGSeC using the implemented system.



Architecture overview.

Adaptation and Evolution (Suda). Adaptation and evolution through natural selection are key biological concepts that we apply in the Bio-Networking Architecture. We have designed evolutionary mechanisms for cyber-entities and demonstrated through simulations that network applications implemented using the Bio-Networking Architecture evolve and adapt to a wide variety of network environments without relying on central control.

The research on adaptation and evolution in the Bio-Networking Architecture has focused on the development of a generic framework where multiple types of cyber-entities interact and co-evolve in the network. We have extended the simulator to accommodate interactions among various types of cyber-entities. The extended simulator has been used to investigate how the cyber-entities evolve their interactions through natural selection.

Design and Implementation of the Bionet Platform (Suda). The middleware platform in the Bio-Networking Architecture, the Bionet platform, provides reusable software components for developing, deploying and executing cyber-entities. The components abstract low-level operating and networking details (e.g., I/O, concurrency, messaging and network connection management), and implement high-level runtime services that cyber-entities use to perform their services. The components in the Bionet platform are based on several biological concepts (e.g., energy exchange and pheromone emission). We have designed and implemented the Bionet platform, and have empirically demonstrated that the Bionet platform is efficient, scalable, reusable, and significantly simplifies development of network applications.

We have worked actively in the Object Management Group (OMG) to reflect the key designs of the Bio-Networking platform in the OMG Super Distributed Objects specification. With partner organizations such as Hitachi, Ltd. and GMD FOKUS (a national institute for computer science in Germany), we have developed a formal specification for Super Distributed Objects, which has been adopted as a standard by the OMG.

Protocols for Distributed Discovery in Peer-to-Peer Systems (Suda). In the Bio-Networking Architecture, a network application is implemented as a collection of distributed cyber-entities. Cyber-entities are distributed over a network and also migrate among network nodes. In such an environment where network objects are distributed over a network and dynamically migrate, there is a need for a protocol that is capable of locating network objects.

To locate distributed network objects, we have developed two distinct discovery protocols. In both discovery protocols, network objects contain relationships (links or pointers) to one another. Discovery queries originate from a network object and travel from network object to network object through relationships in a decentralized manner. Each network object contains a set of keywords that describe the contents of the network object. Relationships of network objects contain information (e.g., keywords) regarding the relationship partner, providing a mechanism to guide discovery queries.

The first discovery protocol introduces keyword similarity among network object to organize them into clusters as well as relationship history to track relationship partner's past discovery performance. The second discovery protocol utilizes the preference of discovery results (discovery hits) by the discovery originator (i.e., a user or a network object), and applies probabilistic forwarding based on the preference of discovery hits to improve the adaptability of discovery to dynamic environments.

We conducted extensive simulations to investigate the adaptability of the two discovery protocols to dynamic and heterogeneous network environments. In the simulated dynamic network, network objects may move, affecting connectivity of network objects and potentially making some network objects temporarily unreachable and unavailable. Moreover, some network objects (e.g., data such as Web pages and files) might continuously change their contents and properties (e.g., keywords describing their contents, and times that they are last updated). In the simulated heterogeneous network, network objects have heterogeneous capabilities (e.g., processing power, storage size) in forwarding discovery queries and users have heterogeneous requirements for discovery (e.g., requiring fast discovery, low overhead discovery). Simulation results have demonstrated the adaptability of the two discovery protocols to dynamic and heterogeneous network environments.

4.7 Middleware for Supporting Application Programs with End-to-End QoS

TAO Pluggable Protocols Framework (Schmidt): One of the problems with conventional middleware solutions is that they hard-code the distributed object programming model together with the use of general-purpose network protocols (such as TCP/IP) that are inadequate to meet the QoS needs of distributed real-time applications. To overcome this problem, another part of our PERC project has focused on enhancing the TAO pluggable protocols framework shown in the following figure so that developers can concentrate on designing their application-specific logic without having to worry about timing, fault tolerance, mobility, or security issues. The key TAO pluggable protocols framework components are described briefly below:

- **ORB Messaging Component:** This component is responsible for implementing ORB messaging protocols, such as the standard CORBA General InterORB Protocol (GIOP) ORB messaging protocol, as well as custom Environmentally-Specific InterORB Protocols (ESIOPs). An ORB messaging protocol defines a data representation, an ORB message format, an ORB transport protocol or transport adapter, and an object addressing format. Within this framework, ORB protocol developers are free to implement optimized Inter-ORB protocols and enhanced transport adaptors, as long as they respect the standard CORBA interfaces.
- **ORB Transport Adapter Component:** This component maps a specific ORB messaging protocol, such as GIOP or DCE-CIOP, onto a specific instance of an underlying transport protocol, such as TCP, SCTP, or ATM. The figure shows an example in which TAO's transport adapter maps the GIOP messaging protocol onto TCP, the standard mapping called IIOP. In this case, the ORB transport adapter combined with TCP corresponds to the Transport layer in the Internet reference model. However, if ORBs communicate over an embedded interconnect, such as a VME bus, the bus driver and DMA controller provide the Transport layer in the communication infrastructure.
- **ORB Policy Control Component:** We have defined QoS APIs that allow applications to specify their QoS requirements using industry-standard CORBA IDL interfaces. In particular, TAO's pluggable protocols framework provides an extensible policy control component that implements the QoS framework defined in the CORBA Messaging and Real-Time CORBA specifications. This component allows applications to control the QoS attributes of configured ORB transport protocols. Example policies for pluggable protocols include buffer pre-allocations, fragmentation, bandwidth reservation, and maximum transport queue sizes. Policies in CORBA can be set at the ORB, thread, or object level. Thus, application developers can set global policies that take effect for any request issued in a particular ORB. Moreover, these global settings can be overridden on a per-thread basis,

a per-object basis, or even before a particular request. CORBA's Policy framework provides fine-grained control over the ORB behavior while providing simplicity in the common case.

QoS-Enabled Network Protocols Supported by TAO's Pluggable Protocols Framework (Schmidt): We have implemented and/or integrated a wide range of network protocols into TAO's pluggable protocols framework. These protocols can be classified into the following main categories:


- **General-Purpose Network Protocols**, such as TCP/IP, SSL over TPC/IP, UDP/IP, UNIX domain sockets, and shared memory
- **Embedded System IUnderconnects**, such as VME, Fibrechannel, and SCRAMNet.
- **QoS-Enabled Network Protocols**, such as IntServ, DiffServ, and the Stream Control Transmission Protocol (SCTP).

Our accomplishments include integrating DiffServ and SCTP support into TAO's pluggable protocols framework. By adding DiffServ capabilities to the TAO ORB and then mapping Real Time CORBA priority values to Diffserv service classes, we have provided a rich middleware platform for developing and deploying QoS-enabled distributed real-time applications that can preserve task priorities via OS schedulers and network prioritization mechanisms end-to-end. Similarly, the Signal Control Transmission Protocol (SCTP) provides a highly configurable, connection-oriented, message and bytestream transport service. SCTP exposes a large set of parameters that can be configured via TAO's QoS policy framework to customize the connection performance to specific application requirements. Properties that can be customized include message ordering semantics (*e.g.*, ordered or unordered), reliability semantics (*e.g.*, retransmit timeouts and max retransmit tries), connection multiplexing (*e.g.*, number of streams), and network path multiplexing (*e.g.*, network interface set). SCTP's customizable properties can be leveraged by distributed real-time CORBA middleware and applications to reduce the complexity of developing high performance, fault tolerant systems.

ZEN (Klefstad and Schmidt): We have designed and developed ZEN, which is a freely available open source middleware that is compliant with most of the features defined in the CORBA 2.3 specification. Its design is based on many of the patterns, techniques, and lessons learned from the development of TAO. ZEN is implemented in Java to maximum ease of use and has a micro-kernel architecture that minimizes footprint for memory constrained distributed real-time embedded systems.

ZEN's architecture (see the figure below) is based on the layered pluggability model, which allows unused ORB components to be factored out of the ORB core to minimize middleware footprint. An extensible component based POA architecture allows customization based on policies, and active de-multiplexing optimizations that associate client requests with target objects in constant time, regardless of the POA hierarchy. Cross-cutting concerns that cannot be captured within a module or a class, such as logging, are implemented in ZEN using Aspect Oriented Programming techniques. The Portable Interceptor specification has been implemented using AspectJ.

RTZen (Klefstad). We have also designed and developed RTZen, the successor of Zen. RTZen is implemented using the Real-time Specification of Java (RTSJ) and is compliant with most of the features defined in the CORBA 2.3 specification and major portions of the Real-Time CORBA specification. Like its predecessor, it is freely available and open source.



Unlike monolithic object request brokers, ZEN's pluggable micro-kernel architecture can be customized based on design preferences, allowing greater flexibility and optimization.

RTZen's design is based on newly discovered RTSJ design patterns, in addition to patterns, techniques, and lessons learned from the development of ZEN. The key innovation behind RTZen is its memory hierarchy, which demonstrates how the RTSJ's scoped memory can provide highly predictable memory management. Besides providing the benefits of ZEN, RTZen offers much higher predictability and throughput due to its scoped memory hierarchy, illustrated in the figure below. The ORB Core of RTZen facilitates the use of key RTSJ features (e.g., immortal memory, scoped memory, and no-heap real-time threads) to enhance middleware predictability. Real-time concurrency and dispatching mechanism in RTZen is provided via thread-pooling and reactive I/O mechanisms. Active de-multiplexing optimizations associates client requests with target objects in constant time, regardless of the POA hierarchy. RTZen also uses intelligent component creation strategies to minimize the time and space overhead for non-real-time applications.


RTZen runs on top of both interpreted and ahead-of-time compiled RTSJ platforms, as well as conventional Java Virtual Machines. Currently it provides real-time predictability on the TimeSys Linux operating system compiled using TimeSys's RI, Purdue's Open Virtual Machine, and Washington University's jRate. Associated with RTZen is a tool, called IsoLeak, that allows the visualization of hierarchies scoped memory areas of RTSJ applications and locates memory leaks.

ZEN-kit: Graphical Customization of Real-Time CORBA for Embedded Systems (Klefstad). Customizing real-time CORBA for an application can considerably reduce the size of the middleware and improve its performance. However, customizing middleware is an error-prone task and requires deep knowledge of the CORBA standard as well as the middleware design. The RT-CORBA features of RTZen are modularized in a hierarchical form (see the figure below) so that they can be customized at different levels of granularity. This customization is achieved through modularizing the middleware so that features may be inserted or removed based on the application requirements.

ZEN-kit is a graphical tool that embodies the principles of modular customization of RTZen for low-footprint devices. ZEN-kit employs two mechanisms for the modularization of RTZen components and aspects. It provides a configuration strategy for the customization of real-time middleware to achieve low-footprint ORBs via ZEN-kit, a graphical tool for composing customized real-time middleware.

Empirical Evaluation of Java-based CORBA Component Model for Embedded and Real-Time Systems (Klefstad). Component technology can overcome many limitations of conventional Object Request Brokers (ORBs) developing distributed, real-time, and embedded Distributed Real-time Embedded (DRE) applications. The CORBA Component Model (CCM) enables the composition and reuse of software components and the configuration of key non-functional aspects of DRE systems such as timing, fault-tolerance, and security. However, CCM can introduce an additional overhead to the runtime performance of middleware. CCM implementations also add significantly to the code size of an application. Hence, the overhead and effects of using CCM need to be evaluated in order to effectively employ it in the design of high-reliability DRE applications.

We empirically evaluated the performance of OpenCCM, a Java-based implementation of the CCM standard. OpenCCM was configured with OpenORB, a desktop Java ORB and ZEN, a real-time Java ORB. We provided throughput and latency measurements of method invocations, as well as measurements for both the ORBs configured with and without OpenCCM. We measured the performance overhead of using OpenCCM on each ORB. We also measured the additional memory requirement introduced by the CCM implementation and provided an analysis of the pros and cons of using component technology in DRE systems



A screen shot of the graphical user interface of ZEN-kit.

Integration of Load Balancing with Fault Tolerance in Distributed Systems (Moser and Melliar-Smith). Existing CORBA load balancing systems lack fault tolerance and CORBA fault tolerance infrastructures do not provide load balancing. Integration of fault tolerance with load balancing provides a better quality of service to the users of the applications by making the infrastructure and applications more robust, available and reliable, and by providing better scalability, response times and throughput.

Aditya Singh, Louise Moser and Michael Melliar-Smith integrated TAO's Load Balancer for CORBA with Eternal System's FTORB fault tolerance infrastructure. TAO's Load Balancer balances the load of the clients across multiple copies of the server using a random assignment policy, while FTORB replicates CORBA applications using either active or passive replication and thus makes them fault-tolerant. The load balancer and the applications are replicated so that the applications can continue to operate despite faults. Overheads are higher than for either load balancing or fault tolerance alone, but are still reasonable.

Protocols and Analysis Model for QoS Support in Networks (Suda). We have investigated various network protocols dealing with QoS under different network environments, from ad-hoc sensor networks composed of capability-limited small devices (i.e., sensors) to the Internet core comprising powerful backbone routers. In the sensor network QoS project, we have developed a QoS analysis model for sensor networks, which analyzes a sensor network holistically. In the Internet QoS project, we have investigated mixed UDP and TCP traffic, as well as inter-domain mechanisms. The issues that we considered include improving the efficiency of best-effort traffic and the management of inter-domain premium traffic. Results of the research effort in each project are described in more detail below.

Sensor Network QoS Analysis Model (Suda) In Wireless Sensor Networks (WSN) Quality of Service (QoS) has been isolated and focused on either certain functional layers or certain application scenarios. We have proposed a holistic wireless sensor network QoS framework as the next step in wireless sensor network QoS research. Using the proposed framework, application designers will be able to unambiguously specify QoS requirements, and system engineers will be able to holistically trace relationships and evaluate tradeoffs in the design and deployment of QoS-aware wireless sensor network applications. The tradeoff relationships can be used to further understand how QoS requirements for wireless sensor networks interact.

We have approached the development of this framework by defining wireless sensor network QoS requirements within a WSN reference architecture, and then analyzing how the defined wireless sensor network QoS requirements impact each other within this QoS framework. We have also proposed a methodology with formal rules to identify such tradeoffs.

Router Queue Management for Mixed Traffic (Suda). We have proposed a robust congestion controller against a mix of unresponsive and responsive traffic. In this project, we have designed a new framework for Internet routers with Active Queue Management (AQM) algorithms to be more robust against a mixture of different traffic sources running different protocols (e.g., FTP, HTTP, UDP, etc.). We are the first to introduce a wavelet-based de-noising technique to remove the impact of unresponsive traffic on the AQM controller. By introducing a separate traffic de-noising filter, the AQM controller is capable of removing the impact of unresponsive traffic on the performance of the responsive traffic (i.e., long-lasting TCP traffic). This technique enables a new paradigm in the design of AQM controller with mixed traffic in the Internet.

We have designed the wavelet de-noising filter that is used to separate and remove unresponsive traffic from the responsive traffic in the Internet traffic. We have conducted extensive simulations to examine the efficiency and robustness of the proposed design.

Inter-domain QoS Automation Using Economics (Suda). We have investigated the inter-domain Resource Exchange (iREX) architecture for the automated deployment of an end-to-end inter-domain QoS policy among multiple Internet Service Providers. iREX uses economics and fully distributed mechanisms to self-manage the deployment of this policy while promoting congestion-avoidance by enabling a distributed resource selection process that selects the least congested inter-domain deployment path.

We have designed the architecture and protocols for iREX. We have investigated the performance of the proposed protocols through extensive simulations. The simulation results show that IREX protocols perform well under different simulation configurations.

| Network Architecture | Application | Socket Layer | Transport | Routing | Traffic Model | Tagger Marker | Buffer Management | NI Scheduling |
|-------------------------|---------------------|----------------|--|--|---|--------------------------------------|---------------------------------------|--|
| Best Effect Services | <i>FTP, FSP WWW</i> | <i>BSD 4.3</i> | <i>TCP-Reno TCP-Tahoe TCP-Vegas TCP Sack UDP</i> | <i>RIP (DV) OSPFv2 Multicast shortest path tree Multicast Steiner tree DVMRP MOSPF CBT</i> | | | <i>Drop-Tail</i> | <i>FIFO</i> |
| Differentiated Services | | | | | | <i>Token Bucket TSW ETSW</i> | <i>RED FRED SRED BRED</i> | <i>FIFO</i> |
| Integrated Services | | | <i>RSVP</i> | <i>Unicast QoS routing QoS-enhanced OSPFv2 QoS-enhanced CBT</i> | <i>Periodic message (CBR) Peak rate model Leaky bucket model Token bucket model IETF/Intserv Flowspec (r,t)-smooth model (C,D)-smooth model</i> | | | <i>RM EDF DCTS VirtualClock LFVC PGPS STFQ WF²Q</i> |

FTP: file transfer protocol

BSD: Berkeley socket distribution

RIP: routing information protocol

DVMRP: distance vector multicast routing protocol

CBT: core based tree protocol

TSW: time sliding window

RED: random early drop

SRED: stable random early drop

RSVP: Resource reservation protocol

RM: rate monotonic

DCTS: distance constrained task system

SCFQ: self-clocked fair queuing

STFQ: start time fair queuing

FSP: file service protocol

TCP: transmission control protocol

OSPF: open shortest path first

MOSPF: multicast extension to OSPF

FIFO: first in first out

ETSW: enhanced time sliding window

FRED: fair random early drop

BRED: balanced random early drop

CBR: constant bit rate

EDF: earliest deadline first

LFVC: leap forward virtual clock

PGPS: packet-by-packet generalized processing sharing

WF²Q: worst-case fair weighted fair queuing

Protocol classes supported in J-Sim.

4.8 Tools for Supporting Network Simulation/Emulation

Simulation of Large Distributed Application-Level Protocols (Baker): In our work with ad hoc networks, as well as other work with peer-to-peer networks, we have found a need for high-level simulation tools, that is, tools that can simulate application-level protocols in large networks with large numbers of flows over long periods of time. Packet-level simulation tools, such as ns-2 and the new JavaSim (work in this project), are designed to provide a level of detail that is not geared, for instance, for simulation of a month's worth of application behavior. We have thus designed and begun implementation of Narses, a new simulator targeted towards large distributed applications. The goal of Narses is to validate the behavior of large applications efficiently using network models of varying levels of detail. For efficiency, our lowest-level construct is not a packet but an application "flow." We introduce several assumptions that allow us to simulate many flows over many nodes over long periods of time. One assumption is that the bandwidth bottlenecks will be the last-hop links and not links internal to the Internet. These assumptions mean that Narses is not appropriate for all topologies, but it is appropriate for all of the topologies we have so far desired to simulate, including large Internet topologies with thousands of nodes and tens of thousands of flows. While fine-grained accuracy of simulated runtimes is not the main goal of Narses, early results with our most detailed network model, as compared to ns-2, show up to a five times speed up in simulation time, with a 53% decrease in memory consumption while maintaining a reasonable degree of accuracy (within 8% on average).

Design, Implementation, and Module Enhancement of J-Sim (Hou): We have carried out extensive research and development tasks to realize an extensible, reusable, and component-based network simulation and emulation environment, as shown in the table and described in more detail below.

Design and Implementation of Autonomous Component Architecture (Hou): We have designed a component-based software architecture, called the Autonomous Component Architecture (ACA), that deploys a message-passing, independent execution model to more closely mimic hardware systems, in terms of how components are specified and assembled and how components interact with one another. To specify, implement, and evaluate ACA, and to explore its application to building large-scale, extensible network simulation/emulation environments, we have carried out the following tasks:

- (1) We have implemented a proof-of-concept version of ACA in Java. Through the lessons learned from implementation and experimentation, we have continued to refine/enrich the implementation. For example, we have designed and implemented an execution context management and scheduling mechanism in the ACA runtime. We have also fine-tuned the performance of the ACA implementation to reduce the component overhead in terms of execution time and memory usage.
- (2) To explore how ACA facilitates the building of large-scale, extensible, and reusable software systems, we have built, on top of the ACA implementation, a compositional network simulation environment, called J-Sim. In summary, we have devised a packet-based network modeling framework, called the Extensible Internetworking Framework (EIF), and have implemented on top of the ACA and EIF implementation, a suite of network protocols in the Internet best-effort service, integrated service, and differentiated service architectures. The table gives the set of protocol classes currently supported in J-Sim. We have also extended J-Sim to include components in mobile wireless environments, i.e., antenna propagation models, terrain models, IEEE 802.11, power saving mode, and ad hoc routing classes (DSR/AODV).

Performance Evaluation (Hou): We have conducted extensive stress tests, compared the performance of J-Sim against ns-2 and SSFNET, and made a detailed qualitative and quantitative comparison (see <http://www.javasim.org/comparison.html>). In spite of the overhead inherited from the component architecture, J-Sim demonstrates better scalability (both in terms of simulation completion time and experiment setup time) in large simulation scenarios (e.g., in the case that the number of nodes ≥ 5000) on a dual-Pentium III-600 MHz, 256MB RAM PC.

Module Enhancement and Extension (Hou): We have extended J-Sim to include components for carrying out simulation for wireless sensor networks. In particular, to enable high-fidelity simulation that reflects how Physical layer characteristics can impact the performance, we have designed and implemented detailed models that characterize (i) Physical layer characteristics such as signal propagation, signal attenuation due

to terrains/foliages, multi-path fading, and signal interference, and (ii) power consumption models in the CPU, memory access, NIC processing, coding/modulation, and other associative circuitry (such as acoustic sound, seismic, or temperature sensors and actuators). We are in the process of incorporating real-life traces that characterize how Physical layer characteristics change with environmental effects (terrain, temperature, obstacles, number of concurrent communication activities).

We have also extended J-Sim to realize network emulation. Specifically, we have developed a complete Java-compliant socket layer on which real applications (e.g., web/ftp servers and audio/video applications) can be readily ported, thus realizing top-down network emulation. We have also leveraged the packet filter facility to intercept real-life packets at the device driver level and redirect them to J-Sim, thus realizing bottom-up network emulation.

Development of Fluid-Model-Based Simulation Techniques for Expediting Simulation (Hou): The major obstacle in packet-level network simulation is the large number of packets that have to be simulated in order to produce accurate results. Each packet generates a number of events (e.g., arrival of a packet at the router, its departure, and its queuing, and buffer depletion) on the path from the source to the destination. In wireless environments where high-fidelity results can be obtained only with simulation at the signal level, the problem is even worse. Due to the broadcast nature of a wireless channel, transmission of a signal must be received and processed by all nodes operating on the same channel (and neighboring channels if co-channel interference is taken into account). This characteristic of a wireless channel implies that one signal transmission event will generate numerous signal receipt events. As the CPU time required is roughly proportional to the number of events that have to be processed, packet-level simulation easily becomes computationally expensive, if not infeasible, when the network size and/or the amount of traffic is extremely large. In our simulation study, we have shown that to simulate a typical WiFi scenario in which n nodes operate in a 802.11-operated wireless LAN, each of which sends CBR traffic at the rate of 0.5 Mbps (with packet size set to 25 bytes), it takes 17083 seconds (4.75 hours) in real time to carry out a 60-second simulation run in the case of $n=100$ nodes.

To solve this problem, we incorporated, for the sake of reducing the number of events, theoretical models in the literature as well as developed by this PERC project, into the simulation. Conceptually, a certain portion of the network traffic (in the MAC layer) is modeled as a continuous fluid flow, rather than as discrete packet instances. A cluster of closely spaced packets is modeled as a single fluid chunk with a fluid rate. An analytical model (in the form of a set of differential equations) is then derived to characterize the system evolution. The simulation environment (with the theoretical models incorporated) keeps track of the fluid rate changes at the traffic sources and router queues. Because a large number of packets is abstracted as a single fluid chunk, less computation is needed to simulate network traffic. We have had success in applying the aforementioned theoretical model-based simulation to congested 802.11 wireless networks and TCP congestion control.

5. Impact

5.1 Protocols for Supporting Real Time Resource Management

- Jennifer Hou has taken to a rigorous level the pioneering work of Tuan and Park on using LRD for congestion control and a simple heuristic estimation scheme, based on conditional expectations, for exploration of the correlation structure. In her work, the traffic prediction is rigorously made with the use of a LMMSE predictor, and the calculation of packet dropping probability in AQM and/or the window adjustment in TCP congestion control is pinpointed in the context of steady-state dynamics.

5.2 Protocols and Infrastructure for Supporting Fault Tolerance

- Michael Melliar-Smith, Louise Moser and their Ph.D. student Wenbing Zhao received the best paper award for their work on performance measurement and analysis of a multicast group communication protocol for a fault tolerance infrastructure.

- Michael Melliar-Smith and Louise Moser participated in the development of the OMG's Fault Tolerant CORBA and Online Software Upgrades standards.
- Michael Melliar-Smith and Louise Moser wrote the OMG's Reliable Ordered Multicast Request for Proposals.
- Michael Melliar-Smith and Louise Moser participated in the development of the Service Availability Forum's Application Interface Specification.
- Michael Melliar-Smith and Louise Moser designed, and supervised the implementation and release of, Eternal Systems' Fault Tolerant CORBA product to NSWC, Dahlgren.
- Michael Melliar-Smith and Louise Moser designed, and supervised the implementation and release of, Eternal Systems' Duration checkpointing and TCP failover product.
- Michael Melliar-Smith and Louise Moser invented technology and prepared patent applications on Online Software Upgrades, High Availability, Fault Tolerance, Message Logging, Checkpointing and Fault Recovery, which were transferred and assigned to Eternal Systems.

5.3 Methodology and Protocols for Supporting Security

- Karl Levitt and his students developed intrusion detection methodology that is suitable for a variety of MANET protocols, including AODV, OLSR and DRCP.
- Karl Levitt and his students developed a formal reasoning and verification methodology applicable to intrusion detection rules in the wired as well as wireless MANET environment.
- Karl Levitt and his students developed automatic response models for application in the wired as well as wireless MANET environment.
- Karl Levitt and his students developed an intrusion detection architecture suitable for bandwidth constrained MANETs.
- Karl Levitt and his students developed worm defense strategies.
- Michael Melliar-Smith, Louise Moser and their Ph.D. student Kim Kihlstrom received the Wilkes Award for their work on Byzantine fault detectors.
- Elizabeth Belding-Royer and her Ph.D. students developed and released simulation code for ARAN, an authenticated routing protocol for ad hoc networks.
- Elizabeth Belding-Royer and her Ph.D. students developed and released implementation code for ARAN that runs on small handheld devices.
- Elizabeth Belding-Royer and her Ph.D. students developed and released AODVSTAT, the first ad hoc intrusion detection system, both as a simulation and implementation in a testbed.

5.4 Protocols for Supporting Scalable Real-Time Multicasts

- Michael Melliar-Smith and Louise Moser invented a low-latency real-time multicast protocol, which their Ph.D. student Wenbing Zhao implemented. The technology was transferred to Eternal Systems.
- The research of Jennifer Hou on QoS-driven multicast routing has direct applicability to the development and implementation of network protocols/solutions in support of resource management in network-centric warfare scenarios. Jointly with Fujitsu Labs of America, Professor Hou has presented the QoS-driven multicast routing protocol in the Inter-Domain Multicast Routing Working Group of IETF.

5.5 Protocols for Supporting Mobile Wireless Environments

- J. J. Garcia-Luna-Aceves developed and released the ROMA medium access control protocol, which was used by Raytheon in demonstrations to DARPA in the FCS program. ROMA was the first MAC protocol based on transmission scheduling applicable to directional antennas.
- J. J. Garcia-Luna-Aceves demonstrated that AODV is not fail-safe when nodes can delete routing state, which may lead to counting-to-infinity and looping, and published several approaches to solve AODV's limitations.
- Jennifer Hou invented the LMST algorithm, the first localized topology control algorithm reported in the literature. Professor Hou has worked with Lockheed Martin to implement the algorithm in their testbed, as part of the DARPA control-based MANET program.
- Jennifer Hou conducted one of the first comprehensive and systematic studies of understanding the performance limits of wireless sensor networks, based on several node and network characteristics (critical power, critical density, critical radius, network coverage, connectivity, capacity, and lifetime).
- Based on the above derived critical conditions, Jennifer Hou and her students designed localized algorithms and protocols that approach performance limits derived in their asymptotic study.
- Leveraging on J-Sim, Jennifer Hou carried out a detailed and high-fidelity simulation study, validating these conditions in practice. The extensions to J-Sim will be released in as open source, as has been done for the previous J-Sim release (<http://www.j-sim.org>).
- Michael Melliar-Smith and Louise Moser and their students developed natural language speech input applications, and demonstrated effective speech recognition on an OQO mobile handheld device. This technology has attracted considerable interest, in particular, from
 - Hossein Moiin, VP, T-Mobile
 - Norman Winarsky, VP, SRI International
 - Robert Scott, Nokiawho and are considering possible commercialization of this technology.

5.6 Bio-Networking Architectures

- BBC News featured an article that overviews the motivation and vision of the Bio-Networking Architecture developed by Tatsuya Suda. See <http://news.bbc.co.uk/1/hi/sci/tech/764085.stm> In this article, adaptive behaviors of network applications in the Bio-Networking Architecture are described.
- Key protocol designs in the Bio-Networking Architecture have been adopted as a reference architecture at the Super Distributed Objects group of the Object Management Group.

5.7 Middleware for Supporting Application Programs with End-to-End QoS

- Raymond Klefstad and Douglas Schmidt developed and released ZEN and RTZen, which implements RT-CORBA in real-time Java.
- Raymond Klefstad developed and refined the pluggable transport protocol framework in RTZen, which was later replaced by the Extensible Transport Framework, which has become a part of the OMG standard.
- Raymond Klefstad released RTZen, the first available open-source RTSJ CORBA middleware implementation
- RTZen is the largest RTSJ software to date and is used as a test program for real-time JVMs, e.g. OVM by Purdue University and jRate.

- RTZen, developed by Raymond Klefstad and his students, was evaluated by JPL, NASA, and Raytheon . They presented RTZen at the “Java day” event at Raytheon . RTZen was considered for use in the DD(X)
- RTZen, developed by Raymond Klefstad and his students, was used in PrismJ, which is a part of a capstone demo for the DARPA PCES (Program Composition for Embedded Systems) program.
- Michael Melliar-Smith, Louise Moser and their Ph.D. student Vana Kalogeraki participated in the development and release of the OMG Dynamic Scheduling standard.

5.8 Tools for Supporting Network Simulation/Emulation

- Jennifer Hou released as open source J-Sim, which provides an extensible, reusable network simulation and emulation environment. Since its first release in October 2001, more than 5000+ users have downloaded the software, including NIST, Oak Ridge National Laboratory, Fujitsu Labs of America, IBM, CMU, University of Toronto, Purdue University, Dartmouth College, Georgia Tech, and Renesys, Inc, many of whom have returned for the latest release, V. 1.3. Extensions to wireless networks, wireless sensor networks, active networks, MPLS, and security mechanisms have been made, some of which are provided by J-Sim users from Europe, Asia, and the States.

6. Education and Student Training

This MURI grant supported the following Postdoctoral, Graduate Student and Undergraduate Student Researchers:

University of California, Santa Barbara, Professor Michael Melliar-Smith, Professor Louise Moser

- Ruppert Koch, Graduate Student Researcher, Ph.D., Fall 2000, Postdoctoral Researcher, 2002-2003
- Prithviraj Dasgupta, Graduate Student Researcher, Ph.D., Spring 2001, Postdoctoral Research, Summer 2001
- Wenbing Zhao, Graduate Student Researcher, Ph.D., Spring 2002, Postdoctoral Researcher, 2002-2004
- Lauren Tewksbury, Graduate Student Researcher, Ph.D., Spring 2001
- Nitya Narasimhan, Graduate Student Researcher, Ph.D., Spring 2001
- Nearchos Paspali, Graduate Student Researcher, M.S., Fall 2002
- Aditya Singh, Graduate Student Researcher, M.S., Spring 2003
- Xiaoli Chen, Graduate Student Researcher, M.S., Summer 2003
- Rachit Chawla, Graduate Student Researcher, M.S., Winter 2004
- Karthik Perumalsamy, Graduate Student Researcher, M.S., Summer 2005
- Brendan Sever, Undergraduate Student Researcher, B.S., Fall 2005
- Rama Alebouyeh, Graduate Student Researcher
- Michael Schuricht, Graduate Student Researcher
- Yi Xin Hu, Graduate Student Researcher
- Shreyas Prasad, Undergraduate Student Researcher
- Zachary Davis, Undergraduate Student Researcher

University of California, Santa Barbara, Professor Elizabeth Belding-Royer, Associate Professor

- Sumit Gwalani, Graduate Student Researcher, M.S., September 2003
- Swaminathan Sundaramurthy, Graduate Student Researcher, M.S., September 2003
- Kavitha Srinivasan, Graduate Student Researcher, M.S., December 2004
- Kimaya Sanzgiri, Graduate Student Researcher,
- Irfan Sheriff, Graduate Student Researcher.

University of California, Irvine, Professor Tatsuya Suda

- Dr. Lubomir Bic, Faculty Member
- Dr. Tadashi Nakano, Postdoctoral Researcher
- Dr. Jun Suzuki, Postdoctoral Researcher
- Ryota Egashira, Graduate Student Researcher
- Akihiro Enomoto, Graduate Student Researcher
- Keita Fujii, Graduate Student Researcher
- Yan Huang, Graduate Student Researcher
- Jun Lu, Graduate Student Researcher
- Michael Moore, Graduate Student Researcher
- Ariffin B. Yahaya, Graduate Student Researcher
- Mei Yang, Graduate Student Researcher
- Takashi Hashimoto, Visiting Researcher
- Dr. Naoto Miyoshi, Visiting Researcher

University of California, Irvine, Professor Raymond Klefstad

- Juan Colmenares, Graduate Student Researcher
- Shruti Gorappa, Graduate Student Researcher
- Trevor Harmon, Graduate Student Researcher, M.S. June 2005
- Jie Hu, Graduate Student Researcher
- Hojjat Jafarpour, Graduate Student Researcher
- Jinhwan Lee, Graduate Student Researcher
- Mark Panahi, Graduate Student Researcher, M.S. December 2004
- Krishna Raman, Graduate Student Researcher, M.S.
- Gunar Schirner, Graduate Student Researcher, M.S.
- Chia-Yen Shih, Graduate Student Researcher
- Yue Zhang, Graduate Student Researcher

University of California, Santa Cruz, Professor J. J. Garcia-Luna-Aceves

- Lichun (Luke) Bao, Ph.D., December 2002
- Soumya Roy, Ph.D., June 2003
- Marc Mosko, Ph.D., May 2004
- Yu Wang, Ph.D., May 2004
- Marco Aurelio Spohn, Ph.D., July 2005
- Ravindra Vaishampayan, Ph.D. Expected 2006
- Hari Rangarajan, Ph.D. Expected 2006

University of California, Davis, Professor Karl. N. Levitt

- Poornima Balasubramanyam, Postdoctoral Researcher
- Jeff Rowe, Postdoctoral Researcher
- Akshay Aggarwal, Graduate Student Researcher, M.S. 2003
- Marcus Tylutki, Graduate Student Researcher, Ph.D. August 2005
- Tao Tsong, Graduate Student Researcher , Ph.D. expected Spring 2006
- Henry Tseng, Graduate Student Researcher
- Angelene Wang, Graduate Student Researcher
- Ivan Balepin, Graduate Student Researcher
- Allen Ting, Graduate Student Researcher

University of Illinois at Urbana Champaign, Professor Jennifer Hou

- Hung-Ying Tyan, Graduate Student Researcher, Ohio State University, Ph.D. Fall 2001
- Yuan Gao, Graduate Student Researcher, Ohio State University, Ph.D. September 2002
- Guanghui He, Graduate Student Researcher, University of Illinois at Urbana Champaign, Ph.D. August 2004
- Ning Li, Graduate Student Researcher, University of Illinois at Urbana Champaign, Ph.D. August 2005
- Honghai Zhang, Graduate Student Researcher, Student, University of Illinois at Urbana Champaign, Ph.D. August 2005

7. Publications During the Period of the Grant

A. Journal Papers

A. Shum, P. M. Melliar-Smith and L. E. Moser, "Design and evaluation of the Fibonacci optical ATM switch," IEEE Transactions on Computers, vol. 50, no. 5 (May 2001), 466-481.

L. E. Moser, P. M. Melliar-Smith, P. Narasimhan, R. R. Koch and K. Berket, "A multicast group communication protocol, engine and bridge for CORBA," Concurrency and Computation: Practice and Experience, volume 13, issue 7 (June 2001), 579-603.

N. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Interceptors for Java Remote Method Invocation," Concurrency and Computation: Practice and Experience volume 13, issue 8-9 (July/August 2001), 755-774.

E. Thomopoulos, L. E. Moser and P. M. Melliar-Smith, "Latency analysis of the Totem single-ring protocol," ACM/IEEE Transactions on Networking, vol. 9, no. 5 (October 2001), 669-680.

K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, "The SecureRing group communication system," ACM Transactions on Information and System Security 4, 4 (November 2001), 371-406.

P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Eternal – A component-based framework for transparent fault-tolerant CORBA," Software Practice and Experience, Theme Issue on Enterprise Frameworks 32, 8 (2001), 771-778.

P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Strongly consistent replication and recovery of fault-tolerant CORBA applications," Computer System Science and Engineering Journal 17, 2 (March 2002), 103-114.

K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, "Unreliable Byzantine fault detectors for solving consensus," Computer Journal 46, 1 (2003), 16-35.

T. Jokiah, F. Herrmann, D. Penkler and L. E. Moser, "The Service Availability Forum Application Interface Specification," Real Time Computing, also in Telecom and Networking (June 2003).

L. E. Moser and P. M. Melliar-Smith, "The Service Availability Forum solution for high availability," CompactPCI Systems, vol. 8, no. 4 (May-June 2004), 36-46.

M. Reitenspiess and L. E. Moser, "The Hardware Platform Interface of the Service Availability Forum," Communications Design (June 2004).

M. Reitenspiess and L. E. Moser, "The Application Interface Specification of the Service Availability Forum," Communications Design (June 2004).

M. Reitenspiess and L. E. Moser, "The Service Availability Forum Application Interface Specification," Boards and Solutions (September 2004), 16-18.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Design and implementation of a consistent time service for fault-tolerant distributed systems," Computer System Science and Engineering Journal, vol. 19, no. 5 (September 2004), 315-323.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Design and implementation of a pluggable fault-tolerant CORBA infrastructure," Cluster Computing: The Journal of Networks, Software Tools and Applications, Special Issue on Dependable Distributed Systems, vol. 7, no. 4 (October 2004), 317-330.

C. Todd and L. E. Moser, "SA Forum specs come of age: New versions based on feedback from implementers," Compact PCI Advanced TCA Systems (January 2005).

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Unification of transactions and replication in three-tier architectures based on CORBA," IEEE Transactions on Dependable and Secure Computing, vol. 2, no. 1 (January-March 2005) 20-33.

P. Dasgupta, L. E. Moser and P. M. Melliar-Smith, "Dynamic pricing for time-limited goods in a supplier-driven electronic marketplace," Journal of Electronic Commerce Research, Special Issue on Dynamic Pricing Policies in Electronic Commerce Springer, vol. 5, no. 2 (April 2005), 267-292.

D. Shahane and L. E. Moser, "COTS becomes highly available," Boards and Solutions (April 2005).

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "End-to-end latency of a fault-tolerant CORBA system," Performance Evaluation Journal, 2005.

K. Sanzgiri, D. LaFlamme, B. Dahill, B. N. Levine, C. Shields and E. M. Belding-Royer, "An authenticated routing protocol for secure ad hoc networks," Journal on Selected Areas in Communications special issue on Wireless Ad hoc Networks, March 2005.

D. C. Schmidt, S. Munjee, S. Flores-Gaitan and A. Gokhale, "Software architectures for reducing priority inversion and non-determinism in real-time Object Request Brokers," Journal of Real-Time Systems, Kluwer, vol. 21, no. 2, 2001.

C. O'Ryan, D. C. Schmidt, F. Kuhns, M. Spivak, J. Parsons Irfan Pyarali and D. L. Levine, "Evaluating policies and mechanisms to support distributed real-time applications with CORBA," Concurrency and Computing: Practice and Experience (Special Issue on Distributed Objects and Applications), Wiley and Sons, vol. 13, no. 2, February, 2001.

D. C. Schmidt, V. Kachroo, Y. Krishnamurthy and F. Kuhns, "Developing next-generation distributed applications with QoS-enabled DPE middleware," IEEE Communications magazine, edited by Abdi Modarressi and Sheshadri Mohan, vol. 17, no. 10, October 2000.

D. C. Schmidt and F. Kuhns, "An overview of the real-time CORBA specification," IEEE Computer, Special Issue on Object-Oriented Real-Time Distributed Computing, edited by Eltefaat Shokri and Philip Sheu, June 2000.

T. Nakano and T. Suda, "Self-organizing network services with evolutionary adaptation," IEEE Transactions on Neural Networks, Special Issue on Adaptive Learning Systems in Communication Networks, vol. 16, no. 5, September 2005.

G. C. Lin, T. Suda and F. Ishizaki, "Loss probability for a finite buffer multiplexer with the M/G/Infinity input Process," Telecommunication Systems, vol. 29, no. 3, July 2005, pp. 183-199.

J. Suzuki and T. Suda, "A middleware platform for a biologically-inspired network architecture supporting autonomous and adaptive applications," IEEE Journal on Selected Areas in Communications, Special Issue on Intelligent Services and Applications in Next Generation Networks, vol. 23, no. 2, February 2005.

Y. Pan, M. Lee, J.B. Kim and T. Suda, "An end-to-end multi-path smooth handoff scheme for stream media," IEEE Journal of Selected Areas of Communications, Special-Issue on All-IP Wireless Networks, vol. 22, no. 4, May 2004, pp. 653-663.

C. Albuquerque, B. Vickers and T. Suda, "Network border patrol: Preventing congestion collapse and promoting fairness in the Internet," IEEE/ACM Transactions on Networking, vol. 12, no. 1, February 2004.

T. Itao, S. Tanaka, T. Suda and T. Aoyama, "A framework for adaptive UbiComp applications based on the Jack-in-the-Net architecture," Kluwer/ACM Wireless Network Journal, vol. 10, no. 3, 2004, pp. 287-299.

F. Ishizaki, G. C. Lin and T. Suda, "An application of the proportional relation to concurrent estimation," IEEE Transactions on Automatic Control, vol. 48, no.10, 2003, pp.1861-1865.

T. Otsuka, M. Matsuo and T. Suda, "Adaptive creation of network application by patchworking application module," Transactions of the Institute of Electronics Communication Engineers of Japan, vol. J86-B, no. 8, 2003, pp. 1404-1416.

T. Itao, T. Nakamura, M. Matsuo, S. Tanaka, T. Suda and T. Aoyama, "Relationships: A mechanism for dynamic and user preference-aware service composition," Information Processing Society of Japan, Journal, vol. 44, no. 3, March 2003.

T. Otsuka, M. Matsuo and T. Suda, "Composition of adaptive service based on patched functional components," Information Technology Letters, vol. 1, no.1, September 2002, pp. 215-216

G. Lin, T. Suda and F. Ishizaki, "Performance analysis of reassembly and multiplexing queuing with long-range-dependent input traffic," Telecommunication Systems, Kluwer Academic Publishers, vol. 20, no.1-2, May-June 2002, pp. 33-58.

T. Suda, M. Matsuo, T. Itao, T. Nakamura, M. Imada, T. Otsuka and S. Tanaka, "Ja-Net adaptive networking architecture for service emergence," Transactions of the Information Processing Society of Japan, June 2002, pp. 616-622.

D. Hong, C. Albuquerque, C. Oliveira and T. Suda, "Evaluating the impact of emerging streaming media applications on TCP/IP performance," IEEE Communications Magazine, April 2001.

T. Suda, T. Itao, T. Nakamura and M. Matsuo, "Adaptive networking architecture for service emergence," Transactions of the Institute of Electronics Communication Engineers of Japan, vol. J84-B, no. 3, 2001, pp. 310-320.

B. J. Vickers, C. Albuquerque and T. Suda, "Source adaptive multi-layered multicast algorithms for real-time distribution," IEEE/ACM Transactions on Networking, vol. 8, no 6, December 2000, pp. 720-733.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Improving route discovery in on-demand routing protocols using two-hop connected dominating sets," Ad Hoc Networks Journal, accepted for publication, 2005.

R. Vaishampayan and J. J. Garcia-Luna-Aceves, "Efficient and robust multicast routing in mobile ad hoc networks," International Journal of Wireless and Mobile Computing, Special Issue on Group Communications in Ad hoc Networks, No. 3, 2005.

J. J. Garcia-Luna-Aceves and S. Roy, "On-demand routing in ad hoc networks using link vectors," IEEE Journal on Selected Areas in Communications, Special Issue on Ad Hoc Networks, vol. 23, no. 3, March 2005.

J. J. Garcia-Luna-Aceves, M. Mosko and Charles Perkins, "A new approach to on-demand loop-free routing in ad hoc networks," accepted for publication in Computer Network, Special Issue on Networking Algorithms, 2004.

Y. Wang and J. J. Garcia-Luna-Aceves, "Modeling of collision avoidance protocols in single-channel multihop wireless networks," ACM Wireless Networks Journal, vol. 10, no. 5, 2004, pp. 495-506.

Y. Wang and J. J. Garcia-Luna-Aceves, "Directional collision avoidance in ad hoc networks," Performance Evaluation, vol. 58, September 2004, pp. 215-241.

Y. Wang and J. J. Garcia-Luna-Aceves, "A new hybrid channel access scheme for ad hoc networks," ACM WINET Journal, Special Issue on Ad-Hoc Networks, vol. 10, no. 4, July 2004.

L. Bao and J. J. Garcia-Luna-Aceves, "Receiver-oriented multiple access in ad hoc networks with directional antennas," Wireless Networks Journal, Special Issue on best papers from Mobicom 2002, vol. 11, 2005, pp. 67-79.

L. Bao and J. J. Garcia-Luna-Aceves, "Distributed dynamic channel access scheduling for ad hoc networks," Journal of Parallel and Distributed Computing, Special Issue on Wireless and Mobile Ad Hoc Networking and Computing, 2002.

H. Y. Tyan, J. Hou and B. Wang, "Many-to-many multicast routing with temporal QoS guarantee," IEEE Transactions on Computers, July 2001.

Y. Gao, J. Hou, and S. Paul, "RACCOOM: A rate-based congestion control scheme for multicasts," IEEE Transactions on Computers, October 2001.

H. Y. Tyan, J. Hou, and B. Wang, "QoS extension to core-based multicast routing," IEEE/ACM Transactions on Networking, July 2002.

G. He, Y. Gao, J. C. Hou, and K. Park, "On exploiting self-similarity of Internet traffic in TCP congestion control, Computer Networks, vol. 45, issue 6, August 2004, 743-766.

W. P. Chen, J. C. Hou, and L. Sha, "Dynamic clustering for acoustic target tracking in wireless sensor networks," IEEE Transactions on Mobile Computing, Special Issue on Mission-Oriented Sensor Networks, vol. 3, no. 3, July-September 2004.

H. Kim and J. C. Hou, "Improving protocol capacity with model-based frame scheduling in IEEE 802.11-operated WLANs," IEEE Journal on Selected Areas in Communications, Special Issue on Design, Implementation and Analysis of Communication Protocols, vol. 22, no. 10, December 2004, pp. 1987-2003.

P. Havinga, J. C. Hou, M. Srivastava and F. Zhao, "Wireless sensor networks," IEEE Wireless Communications Magazine, Special Issue on Wireless Sensor Networks, December 2004.

H. Zhang and J. C. Hou, "Maintaining sensing coverage and connectivity in large sensor networks," Wireless Ad Hoc and Sensor Networks: An International Journal, vol. 1, no. 1-2, January 2005, pp. 89-123.

N. Li, J. C. Hou and L. Sha, "Design and analysis of a MST-based distributed topology control algorithm for wireless ad-hoc networks," IEEE Transactions on Wireless Communications, vol. 4, no. 3, May 2005, pp. 1195-1207.

H. Lim, J. C. Hou, and C. H. Choi, "Provisioning of network distance service on the Internet," ACM/IEEE Transactions on Networking, vol. 13, no. 3, June 2005, pp. 513-526.

H. Zhang and J. C. Hou, "On deriving the upper bound of α -lifetime for large sensor networks," ACM Transactions on Sensor Networks, vol. 1, no. 6, December 2005.

C. Hu and J. C. Hou, "A link-indexed statistical traffic prediction approach to improving IEEE 802.11 PSM," ACM/Kluwer Ad Hoc Networks Journal, Special Issue on Data Communication and Topology Control in Ad Hoc Networks," April 2004.

C. Ko, P. Brutch, J. Rowe, G. Tsafnat and K. N. Levitt, "System health and intrusion monitoring using a hierarchy of constraints," Recent Advances in Intrusion Detection, Lecture Notes in Computer Science, W. Lee, L. Me and A. Wespi, (eds.), vol. 2212, 2001, pp. 190-203.

B. Conference Papers

N. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Interception in the Aroma system," Proceedings of the Java Grande Conference, San Francisco, CA (June 2000), 107-115.

R. R. Koch, L. E. Moser and P. M. Melliar-Smith, "A reliable many-to-many multicast protocol for group communication over ATM networks," Proceedings of the IEEE International Conference on Dependable Systems and Networks, New York, NY (June 2000), 337-346.

P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Transparent fault tolerance for enterprise applications," Proceedings of the International Conference on Advances in Infrastructure for Electronic Business, Science and Education on the Internet L'Aquila, Rome, Italy (July/August 2000).

V. Kalogeraki, P. M. Melliar- Smith and L. E. Moser, "Managing object groups in fault-tolerant distributed object systems," Proceedings of the 13th International Conference on Parallel and Distributed Computing Systems, Las Vegas, NV (August 2000), 509-516.

N. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Transparent consistent replication of JavaRMI objects," Proceedings of the IEEE Conference on Distributed Objects and Applications, Antwerp, Belgium (September 2000), 17-26.

P. Dasgupta, L. E. Moser and P. M. Melliar-Smith, "The security architecture for MAgNET: A mobile agent e-commerce system," Proceedings of the 3rd International Conference on Telecommunications and Electronic Commerce, Dallas, TX (November 2000), 289-298.

P. M. Melliar-Smith, L. E. Moser and V. Kalogeraki, "Dynamic scheduling of distributed method invocations," Proceedings of the IEEE 21st Real-Time Systems Symposium, Orlando, FL (November 2000), 57-66.

V. Kalogeraki, P. M. Melliar-Smith and L. E. Moser, "Dynamic migration algorithms for distributed object systems," Proceedings of the IEEE International Conference on Distributed Computing Systems Phoenix, AZ (April 2001), 119-126.

K. Berket, D. A. Agarwal, P. M. Melliar-Smith and L. E. Moser, "Overview of the InterGroup protocols," Proceedings of the 2001 International Conference on Computational Science, Lecture Notes in Computer Science 2073, Springer Verlag, San Francisco, CA (May 2001), 316-325.

L. A. Tewksbury, L. E. Moser and P. M. Melliar-Smith, "Automatically-generated state transfer and conversion code to facilitate software upgrades," Proceedings of Marcon 2001, Gatlinburg, TN (May 2001).

N. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Transparent failover for Java Remote Method Invocation," Proceedings of Marcon 2001, Gatlinburg, TN (May 2001).

W. Zhao, P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Experimental evaluation of a fault-tolerant CORBA system," Proceedings of the 2001 International Conference on Parallel and Distributed Processing and Applications, Las Vegas, NV (June 2001), 390-396.

E. Royer, P. M. Melliar-Smith and L. E. Moser, "An analysis of the optimum node density for ad hoc mobile networks," Proceedings of the IEEE International Conference on Communications, Helsinki, Finland (June 2001), 857-861.

P. Dasgupta, L. E. Moser and P. M. Melliar-Smith, "Dynamic tiered pricing in a multi-agent economy," Proceedings of the Second International Conference on Internet Computing, Las Vegas, NV (June 2001), 1142-1148.

P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "State synchronization and recovery for strongly consistent replicated CORBA objects," Proceedings of the IEEE 2001 International Conference on Dependable Systems and Networks, Goteborg, Sweden (June/July 2001), 261-270.

K. Kihlstrom, N. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "A secure auction service," Proceedings of the 13th IASTED International Conference on Parallel and Distributed Computing and Systems, Anaheim, CA (August 2001), 599-604.

L. A. Tewksbury, L. E. Moser and P. M. Melliar-Smith, "Live upgrade techniques for CORBA applications," Proceedings of the Third IFIP WG6.1 International Working Conference on Distributed Applications and Interoperable Systems Krakow, Poland (September 2001), 257-271.

L. A. Tewksbury, L. E. Moser and P. M. Melliar-Smith, "Coordinating the simultaneous upgrade of multiple CORBA application objects," Proceedings of the IEEE 3rd International Symposium on Distributed Objects and Applications, Rome, Italy (November 2001), 310-320.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Increasing the reliability of three-tier applications," Proceedings of the IEEE 12th International Symposium on Reliability Engineering, Hong Kong (November 2001), 138-147.

L. A. Tewksbury, L. E. Moser and P. M. Melliar-Smith, "Live upgrades of CORBA applications using object replication," Proceedings of the International Conference on Software Maintenance, Florence, Italy (November 2001), 488-497.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Integrating fault-tolerant CORBA and the CORBA Object Transaction Service," Proceedings of the International Conference on Software and Systems Engineering and Their Applications, Paris, France (December 2001).

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Design and implementation of a pluggable fault-tolerant CORBA infrastructure," Proceedings of the International Parallel and Distributed Processing Symposium, Fort Lauderdale, FL (April 2002).

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "End-to-end latency of a fault-tolerant CORBA infrastructure," Proceedings of the IEEE 5th International Symposium on Object-Oriented and Real-Time Distributed Computing, Washington, DC (April-May 2002), 189-198.

R. R. Koch, L. E. Moser and P. M. Melliar-Smith, "The Totem redundant ring protocol," Proceedings of the IEEE International Conference on Distributed Computing Systems, Vienna, Austria (July 2002), 598-607.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Unification of replication and transaction processing in three-tier architectures," Proceedings of the IEEE 3rd International Conference on Distributed Computing Systems, Vienna, Austria (July 2002), 290-297.

P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Lessons learned in building a fault-tolerant CORBA system," Proceedings of the IEEE International Conference on Dependable Networks and Systems, Washington, DC (June 2002), 39-44.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "End-to-end latency analysis and evaluation of a fault-tolerant CORBA infrastructure," Proceedings of the 2002 International Symposium on Performance Evaluation of Computer and Telecommunication Systems, San Diego, CA (July 2002), 821-830.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "On bootstrapping replicated CORBA Applications," Proceedings of the IEEE 26th Annual International Software and Applications Conference, Oxford, England (August 2002), 239-245.

L. E. Moser and P. M. Melliar-Smith, "Demonstration of fault tolerance for CORBA applications," Proceedings of the DARPA Information and Survivability Conference and Exposition, vol. 2, IEEE Computer Society (2003), 87-89.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Design and implementation of a consistent time service for fault-tolerant distributed systems," Proceedings of the IEEE 4th International Conference on Dependable Systems and Networks, San Francisco, CA (June 2003), 341-350.

R. Koch, S. Hortikar, L. E. Moser and P. M. Melliar-Smith, "Transparent TCP connection failover," Proceedings of the IEEE 2003 International Conference on Dependable Systems and Networks, San Francisco, CA (June 2003), 383-392.

A. Singh, L. E. Moser and P. M. Melliar-Smith, "Integration of fault tolerance and load balancing in distributed systems based on CORBA," Proceedings of the European Dependable Computing Conference, Budapest, Hungary (April 2005), 154-166.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "A reservation-based coordination protocol for business activities," Proceedings of the IEEE International Conference on Web Services, Orlando, FL (July 2005), 49-56.

S. Gwalani, K. Srinivasan, G. Vigna, E. M. Belding-Royer and R. Kemmerer, "An intrusion detection tool for AODV-based ad hoc wireless networks," Proceedings of the Annual Computer Security Applications Conference, Tucson, AZ, December 2004.

K. Sanzgiri, B. Dahill, B. N. Levine, C. Shields and E. M. Belding-Royer, "A secure routing protocol for ad hoc networks," Proceedings of the International Conference on Network Protocols, Paris, France, November 2002.

J. Parsons, I. Pyarali, D. C. Schmidt and R. Cytron, "Achieving end-to-end predictability of the TAO real-time CORBA ORB," Proceedings of the 8th IEEE Real-Time Technology and Applications Symposium, San Jose, CA, September 2002.

R. Klefstad, D. C. Schmidt and C.s O'Ryan, "Towards highly configurable real-time Object Request Brokers," IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC), Washington, D.C., April - May 2002.

C. O'Ryan, D. C. Schmidt, F. Kuhns, M. Spivak, J. Parsons, I. Pyarali and D. L. Levine, "Evaluating policies and mechanisms for supporting embedded, real-time applications with CORBA 3.0," Proceedings of the Sixth IEEE Real-Time Technology and Applications Symposium, Washington ,D.C., May - June 2000.

K. Raman, Y. Zhang, M. Panahi, J. A. Colmenares and R. Klefstad, "RTZen: Highly predictable, real-time Java middleware for distributed and embedded Systems," Proceedings of the ACM/IFIP/USENIX 6th International Middleware Conference, November-December 2005, Grenoble, France.

Y. Zhang, M. Lai, R. Klefstad, R. Lavanya and R. Jayakrishnan, "A distributed, scalable, and synchronized framework for large-scale microscopic traffic simulation," Proceedings of the 8th International IEEE Conference on Intelligent Transportation Systems, September 2005, Vienna, Austria.

K. Raman, Y. Zhang, M. Panahi, J. A. Colmenares and R. Klefstad, "Patterns and tools for achieving predictability and performance with real-time Java," Proceedgins of the 11th IEEE International Conference on Real-Time and Embedded Computing Systems and Applications, 2005, Hong Kong, China.

T. Harmon and R. Klefstad, "VADRE: A visual approach to performance analysis of distributed, real-time systems," Proceedings of the 2005 International Conference on Modeling, Simulation and Visualization Methods, June 2005, Las Vegas, NV.

S. Gorappa, J. A. Colmenares, H. Jafarpour and R. Klefstad, "Tool-based configuration of real-time CORBA middleware for embedded systems," Proceedgins of the 8th IEEE International Symposium on Object-Oriented Real-Time Distributed Computing, Seattle, WA, May 2005

D. Billsus, R. Klefstad and M. J. Pazzani, "Adaptive personalization for mobile content delivery," Proceedings of the 3rd International Conference on Universal Access in Human-Computer Interaction, 2005.

S. Gorappa and R. Klefstad "Empirical evaluation of OpenCCM on CORBA-based Java middleware," Proceedgins of the ACM Symposium on Applied Computing, 2004.

G. Schirner, T. Harmon and R Klefstad, "Late demarshalling: A technique for efficient multi-language middleware for embedded systems," Proceedgins of the Distributed Objects and Applications Conference, 2004.

A. Krishna, D. C. Schmidt and R. Klefstad "Enhancing real-time CORBA via real-time Java features," Proceedings of the International Conference on Distributed Computing Systems, 2004.

N. Shankar and R. Klefstad, "A CORBA framework for distributed service location and creation," Proceedings of the IEEE/ISIJ Symposium on Applications and the Internet, 2004.

A. Krishna, D. C. Schmidt, K. Raman and R. Klefstad, "Optimizing the ORB core to enhance real-time CORBA predictability," Proceedings of the Distributed Objects and Applications Conference, 2003.

A. Krishna, R. Klefstad and Douglas C. Schmidt, "Towards predictable real-time Java Object Request Brokers," Proceedings of the IEEE Real-time Applications and Systems Conference, 2003, pp. 49-57.

R. Klefstad, S. Rao and D. C. Schmidt, "Design and performance of a dynamically configurable, messaging protocols framework for real-time CORBA," Proceedings of the Distributed Object and Component-based Software Systems part of the Software Technology Track at the 36th Annual Hawaii International Conference on System Sciences, January 2003, Big Island of Hawaii, pp. 320-329.

R. Klefstad, A. S. Krishna and D. C. Schmidt, "Design and performance of a modular portable object adapter for distributed, real-time, and embedded CORBA applications," Proceedings of the Distributed Objects and Applications Conference, 2002, pp. 549-562.

A. Corsaro, D. C. Schmidt, R. Klefstad and C. O'Ryan, "Virtual component: A design pattern for memory-constrained embedded applications," Proceedings of the Programming Languages of Patterns Conference, 2002, pp. 4-16.

R. Klefstad, D. C. Schmidt, and C. O'Ryan, "Towards highly configurable real-time Object Request Brokers," Proceedings of the IEEE International Symposium on Object-Oriented Real-Time Distributed Computing, 2002, pp. 437-447.

A. Yahaya and T. Suda, "iREX: Inter-domain resource exchange architecture," IEEE Infocom 2005.

R. Egashira, A. Enomoto, T. Suda, H. Sasaki and H. Iwasaki, "Distributed service discovery using preference," Proceedings of the IEEE/ICST CollaborateCom, December 2005.

J. Lu, L. Bao and T. Suda, "Coverage aware sensor engagement in dense sensor networks," Proceedings of the IFIP International Conference on Embedded and Ubiquitous Computing, December 2005.

J. Lu, L. Bao and T. Suda, "Probabilistic self-scheduling for coverage configuration in sensor networks," Proceedings of the IEEE International Conference on Sensing Technology, November 2005.

M. Yang, Y. Huang, J. Kim, M. Lee, T. Suda and D. Matsubara, "An end-to-end QoS framework with on-demand bandwidth reconfiguration," Proceedings of IEEE Infocom, November, 2004.

K. Fujii and T. Suda, "Dynamic service composition using semantic information," Proceedings of the Second ACM International Conference on Service Oriented Computing, November 2004.

T. Nakano and T. Suda, "Adaptive and evolvable network services," Proceedings of the ACM Genetic and Evolutionary Computation Conference, International Society of Genetic and Evolutionary Computation, Lecture Notes in Computer Science, Springer, vol. 3102, June 2004, pp. 151-162 June 2004.

J. Suzuki and T. Suda, "Design and implementation of a scalable infrastructure for autonomous adaptive agents," Proceedings of the 15th IASTED International Conference on Parallel and Distributed Computing and Systems, November 2003

J. Song, S. Kim, M. Lee, H. Lee and T. Suda, "Adaptive load distribution over multipath in MPLS Networks," Proceedings of the IEEE International Conference on Communications, May 2003.

T. Otsuka, M. Matsuo and T. Suda, "Evolutionary service composition based on patched functional components," Proceedings of the IEICEJ and ISPJ Forum on Information Technology, LM-2, 2002, pp. 215-216.

T. Itao, T. Nakamura, M. Matsuo, T. Suda and T. Aoyama, "Adaptive creation of network applications in the Jack-In-The-Net architecture," Proceedings of the IFIP Networking, 2002.

T. Itao, T. Nakamura, M. Matsuo, T. Suda and T. Aoyama, "Service emergence based on relationship among self-organizing entities," Proceedings of the IEEE/IPSJ Symposium on Applications and the Internet, 2002.

M. Imada, Y. Katayama, M. Matsuo and T. Suda, "A method of service composition using service attributes," Proceedings of the IEICEJ Assurance Symposium, June 2001.

T. Ozaki, J. B. Kim and T. Suda, "Bandwidth-efficient multicast routing protocol for ad-hoc networks," Proceedings of IEEE Infocom, 2001.

T. Itao, T. Suda and T. Aoyama, "Jack-in-the-Net: Adaptive networking architecture for service emergence," Proceedings of the Asian-Pacific Conference on Communications, 2001.

T. Iizuka, A. Lau and T. Suda, "A design of local resource access control for mobile agent in PDA," Proceedings of the Asian-Pacific Conference on Communications, 2001.

S. Ahn, M. Lee and T. Suda, "Topology reconfiguration of an IP network embedded over an ATM network," Proceedings of the Asian-Pacific Conference on Communications, 2001.

S. Song and T. Suda, "Security on energy level in bio-networking architecture," Proceedings of the 3rd International Conference on Advanced Communications Technology, 2001.

M. Wang and T. Suda, "The bio-networking architecture: A biologically inspired approach to the design of scalable, adaptive, and survivable/available network applications," Proceedings of the IEEE SAINT Symposium, 2001.

C. Albuquerque, B. Vickers and T. Suda, "Network Border Patrol," Proceedings of IEEE Infocom, March 2000.

R. Vaishampayan, J. J. Garcia-Luna-Aceves and K. Obraczka, "An adaptive redundancy protocol for mesh based multicasting," Proceedings of the 2005 International Symposium on Performance Evaluation of Computer and Telecommunication Systems, Philadelphia, PA, July 2005.

H. Rangarajan and J. J. Garcia-Luna-Aceves, "Making on-demand routing protocols based on destination sequence numbers robust," Proceedings of the IEEE International Conference on Communications, May 2005, Seoul, Korea.

H. Rangarajan and J. J. Garcia-Luna-Aceves, "Efficient use of route requests for loop-free on-demand routing in ad hoc networks," Proceedings of the IFIP Networking Conference, Waterloo Ontario, Canada, May 2005.

H. Rangarajan and J. J. Garcia-Luna-Aceves, "Path verification for robust instantaneous loop-free routing in ad hoc networks," Proceedings of the IEEE Wireless Communications and Networking Conference, March 2005, New Orleans, LA.

M. A. Spohn and J. J. Garcia-Luna-Aceves, Improving the efficiency and reliability of the route discovery process in on-demand routing protocols," Proceedings of the IEEE Wireless Communications and Networking Conference, March 2005, New Orleans, LA.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "A solution for the location problem in arbitrary computer networks using generic dominating sets," Proceedings of ACM 2005: The 20th Annual ACM Symposium on Applied Computing, Santa Fe, NM, March 2005.

J. J. Garcia-Luna-Aceves and H. Rangarajan, "A new framework for loop-free on-demand routing using destination sequence numbers," Proceedings of the First IEEE International Conference on Mobile Ad-hoc and Sensor Systems, Fort Lauderdale, FL, October 2004.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Improving broadcast operations in ad hoc networks using two-hop connected dominating sets," Proceedings of the First IEEE International Conference on Mobile Ad-hoc and Sensor Systems, Fort Lauderdale, FL, October 2004.

R. Vaishampayan and J. J. Garcia-Luna-Aceves, "Efficient and robust multicast routing in mobile ad hoc networks," Proceedings of the First IEEE International Conference on Mobile Ad Hoc and Sensor Systems, Fort Lauderdale, FL, October 2004.

J. J. Garcia-Luna-Aceves and S. Roy, "On-demand loop-free routing with link vectors," Proceedings of the IEEE International Conference on Network Protocols, Berlin, Germany, October 2004.

H. Rangarajan and J. J. Garcia-Luna-Aceves, "Achieving loop-free incremental routing in ad hoc networks," Proceedings of the 9th IEEE Symposium on Computers and Communications, Alexandria, Egypt, June/July 2004.

H. Rangarajan and J. J. Garcia-Luna-Aceves, "Reliable data delivery in event-driven wireless sensor networks," Proceedings of the 9th IEEE Symposium on Computers and Communications, Alexandria, Egypt, June/July 2004.

H. Rangarajan and J. J. Garcia-Luna-Aceves, "Using labeled paths for loop-free on-demand routing in ad hoc networks," Proceedings of ACM MobiHoc, Tokyo, Japan, May 2004.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Enhancing the route discovery process of on-demand routing in networks with directional antennas," Proceedings of Networking 2004: The Third IFIP-TC6 Networking Conference, Athens, Greece, May 2004.

Y. Wang and J. J. Garcia-Luna-Aceves, "Performance of directional collision avoidance in ad hoc networks," Proceedings of Networking 2004: The Third IFIP-TC6 Networking Conference, Athens, Greece, May 2004.

R. Vaishampayan and J. J. Garcia-Luna-Aceves, "Robust tree-based multicasting in ad-hoc networks," Proceedings of the Workshop on Multihop Wireless Networks, Phoenix, AZ, April 2004.

M. Mosko and J. J. Garcia-Luna-Aceves, "Loop-free routing using a dense label set in wireless networks," Proceedings of IEEE 24th International Conference on Distributed Computing Systems, Tokyo, Japan, March 2004.

M. Mosko, J. J. Garcia-Luna-Aceves and C. Perkins, "Distribution of route requests using dominating-set neighbor elimination in an on-demand routing protocol," Proceedings of IEEE Globecom 2003 - Wireless Communications, San Francisco, CA, December 2003.

Y. Wang and J. J. Garcia-Luna-Aceves, "Broadcast traffic in ad hoc networks with directional antennas," Proceedings of IEEE Globecom 2003 - Wireless Communications, San Francisco, CA, December 2003.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Enhanced dominant pruning applied to the route discovery process of on-demand routing protocols," Proceedings of the IEEE Twelfth International Conference on Computer Communications and Networks, Dallas, TX, October 2003.

Y. Wang and J. J. Garcia-Luna-Aceves, "Channel sharing of competing flows in ad hoc networks," Proceedings of the IEEE Symposium on Computers and Communications, Kemer-Antalya, Turkey, June/July 2003.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Exploiting relative addressing and virtual overlays in ad hoc networks with bandwidth and processing constraints," Proceedings of the 2003 International Conference on Wireless Networks, Las Vegas, NV, June 2003.

L. Bao and J. J. Garcia-Luna-Aceves, "Topology control in ad hoc networks," Proceedings of ACM MobiHoc, Annapolis, MD, June 2003.

Y. Wang and J. J. Garcia-Luna-Aceves, "Collision avoidance in single-channel ad hoc networks using directional antennas," Proceedings of the 23rd IEEE International Conference on Distributed Computing Systems, May 2003, Providence, RI.

Y. Wang and J. J. Garcia-Luna-Aceves, "Throughput and fairness in a hybrid channel access scheme for ad hoc networks," Proceedings of IEEE WCNC 2003, New Orleans, LA, March 2003.

S. Roy and J. J. Garcia-Luna-Aceves, "Node-centric hybrid routing for ad-hoc wireless extensions of the Internet," Proceedings of the IEEE Symposium on Ad Hoc Wireless Networks, Taipei, Taiwan, November 2002.

Y. Wang and J. J. Garcia-Luna-Aceves, "Spatial reuse and collision avoidance in ad hoc networks with directional antennas," Proceedings of IEEE Symposium on Ad Hoc Wireless Networks, Taipei, Taiwan, November 2002.

L. Bao and J. J. Garcia-Luna-Aceves, "Hybrid channel access scheduling in ad hoc networks," Proceedings of IEEE 10th International Conference on Network Protocols, Paris, France, November 2002.

Y. Wang and J. J. Garcia-Luna-Aceves, "Performance of collision avoidance protocols in single-channel ad hoc networks," Proceedings of IEEE 10th International Conference on Network Protocols, Paris, France, November 2002.

S. Roy and J. J. Garcia-Luna-Aceves, "An efficient path selection algorithm for on-demand link-state hop-by-hop routing," Proceedings of IEEE International Conference on Computer Communication and Networks, Miami, FL, October 2002.

L. Bao and J. J. Garcia-Luna-Aceves, "Transmission scheduling in ad hoc networks with directional antennas," Proceedings of ACM Eight Annual International Conference on Mobile Computing and Networking, Atlanta, GA, September 2002.

L. Bao and J. J. Garcia-Luna-Aceves, "Distributed transmission scheduling using code-division channelization," Proceedings of Networking 2002, sponsored by the IFIP Technical Committee on Communication Systems (TC6) working groups Network and Internetwork Architectures (WG 6.2), Performance of Communication Systems (WG 6.3), and Wireless Communications (WG 6.8), Pisa, Italy, May 2002. Reprinted in Lecture Notes in Computer Science 2345, 2002.

B. Wang, H. Y. Tyan, and J. Hou, "QoS-based multicast routing for distributing layered video to heterogeneous receivers in rate-based networks," Proceedings of the IEEE 20th INFOCOM, Tel Aviv, Israel, March 2000.

Y. Gao, Y. Ge, and J. Hou, "Reliable multicasts for core-based multicast routing," Proceedings of the IEEE 8th International Conference on Network Protocols, Osaka, Japan, November 2000.

H. Y. Tyan and J. Hou, "JavaSim: A component-based compositional network simulation environment," Proceedings of the Western Simulation Multiconference -- Communication Networks and Distributed Systems Modeling and Simulation, January 2001.

Y. Gao, J. Hou, and S. Paul, "RACCOOM: A rate-based congestion control scheme for multicasts," Proceedings of the SPIE Conference on Scalability and Traffic Control in IP Networks, August 2001

B. Wang and J. Hou, "An efficient QoS routing algorithm for quorumcast communication," Proceedings of the IEEE 9th International Conference on Network Protocols, November 2001.

H. Y. Tyan and J. Hou, "Design, realization, and evaluation of a component-based, compositional network simulation environment," Proceedings of the Western Multiconference - Communication Networks and Distributed Systems Modeling and Simulation, January 2002.

Y. Gao, G. He, and J. Hou, "On leveraging traffic predictability in active queue management," Proceedings of IEEE INFOCOM 2002, New York, NY, June 2002.

W. P. Chen and J. Hou, "Dynamic, ad-hoc source routing with connection-aware link-state exchange and differentiation," Proceedings of IEEE Globecom, Taipei, Taiwan, November 2002.

G. He, Y. Gao, and J. Hou, "A case for exploiting self-similarity of Internet traffic in TCP congestion control," Proceedings of the IEEE International Conference on Network Protocols, Paris, France, November 2002.

N. Li, J. C. Hou, and L. Sha, "Design and analysis of a MST-based distributed topology control algorithm for wireless ad-hoc networks," Proceedings of IEEE INFOCOM, San Francisco, CA, April 2003

G. He and J. C. Hou, "On exploiting long range dependency in measuring cross traffic," Proceedings of IEEE INFOCOM, San Francisco, CA, April 2003.

Y. Gao and J. C. Hou, "A state feedback control approach to stabilizing AQM queues for ECN-enabled TCP connections," Proceedings of IEEE INFOCOM, San Francisco, CA, April 2003.

L. Y. Zhang, Y. Ge and J. C. Hou, "Energy-efficient real-time scheduling in IEEE 802.11 wireless LANs," Proceedings of IEEE International Conference on Distributed Computing Systems, Providence, RI, May 2003.

Z. Li, P. Zhou, and J. C. Hou, "A fragmentation-based D-MAC protocol for MANETs," Proceedings of IEEE International Conference on Distributed Computing Systems, Providence, RI, May 2003.

R. Zheng, J. C. Hou and L. Sha, "Asynchronous wakeup for ad hoc networks," Proceedings of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing, Annapolis, MD, June 2003.

H. Kim and J. C. Hou, "Improving protocol capacity with model-based frame scheduling in IEEE 802.11-operated WLANs," Proceedings of the ACM International Conference on Mobile Computing and Networking, San Diego, CA, September 2003.

H. Lim, J. C. Hou and C. H. Choi, "Constructing an Internet coordinate system based on delay measurement," Proceedings of the ACM Internet Measurement Conference, Miami, FL, October 2003.

W. P. Chen, J. C. Hou and L. Sha, "Dynamic clustering for acoustic target tracking in wireless sensor networks," Proceedings of the IEEE 11th International Conference on Network Protocols, Atlanta, GA, November 2003.

R. Zheng, J. C. Hou, and L. Sha, "Performance of the IEEE 802.11 power saving mode," Proceedings of the SCS Western Multiconference on Computer Simulation – Communication Networks and Distributed Systems Modeling and Simulation Conference, San Diego, CA, January 2004.

N. Li and J. C. Hou, "Topology control in heterogeneous wireless networks: problems and solutions," Proceedings of IEEE INFOCOM, Hong Kong, China, March 2004.

H. Kim and J. C. Hou, "Network calculus based simulation: Theorems, implementation, and evaluation," Proceedings of IEEE INFOCOM, Hong Kong, China, March 2004.

C. Hu and J. C. Hou, "LISP: An link-indexed statistical traffic prediction approach to improving IEEE 802.11 PSM," Proceedings of IEEE Int'l Conf. on Distributed Computing Systems, Tokyo, Japan, March 2004.

H. Kim and C. J. Hou, "A fast simulation framework for IEEE 802.11-operated wireless LANs," Proceedings of ACM SIGMETRICS, New York, NY, May 2004.

H. Zhang and C. J. Hou, "On deriving the upper bound of alpha-lifetime for large sensor networks," Proceedings of ACM MobiHoc, Tokyo, Japan, June 2004.

A. Sobeih, J. C. Hou, and M. Viswanathan, "Check and simulate: a case for incorporating model checking in network simulation." Proceedings of the Second ACM-IEEE International Conference on Formal Methods and Models for Codesign, San Diego, CA, June 2004.

N. Li and J. C. Hou, "FLSS: A fault-tolerant topology control algorithm for wireless networks," Proceedings of the ACM International Conference on Mobile Computing and Networking, Philadelphia, PA, September 2004.

N. Li and J. C. Hou, "BLMST: A decentralized, power-efficient broadcast algorithm for wireless sensor networks," Proceedings of the IEEE 1st International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks, Dallas, TX, October 2004.

A. Sobeih, B. Yurcik and J. C. Hou, "VRing: A case for building application-layer multicast rings (rather than trees)," Proceedings of the 12th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, Volendam, The Netherlands, October 2004.

S. Kandula, J. C. Hou and L. Sha, "The case for resource heterogeneity in large sensor networks," Proceedings of IEEE Milcom, Monterey, CA, November 2004.

R. Zheng, J. C. Hou and L. Sha, "On time-out driven power management policies in wireless networks," Proceedings of IEEE Global Telecommunications Conference, Dallas TX, November/December 2004.

H. Zhang and J. C. Hou, "On the critical total power for asymptotic k-connectivity in wireless networks," Proceedings of IEEE INFOCOM, Miami, FL, March 2005.

H. Zhang and J. C. Hou, "Capacity of wireless ad-hoc networks under ultra wide band with power constraints," Proceedings of IEEE INFOCOM, Miami, FL, March 2005.

H. Lim and J. C. Hou, "Localization for anisotropic sensor networks," Proceedings of IEEE INFOCOM, Miami, FL, March 2005.

A. Sobeih, W. P. Chen, J. C. Hou, L. C. Kung, N. Li, H. Lim and H. Zhang, "J-Sim: A simulation and emulation environment for wireless sensor networks," Proceedings of the 38th Annual Simulation Symposium, April 2005, San Diego, CA.

G. He and J. C. Hou, "An in-depth, analytical study of sampling techniques for self-similar Internet traffic," Proceedings of the IEEE International Conference on Distributed Computing Systems, June 2005.

C. Hu and J. C. Hou, "A reactive channel model for expediting wireless network simulation," Proceedings of ACM SIGMETRICS, June 2005.

N. Li and J. C. Hou, "Improving connectivity of wireless ad-hoc networks," Proceedings of ACM 2nd International Conference on Mobile and Ubiquitous Systems: Networking and Services, July 2005.

H. Kim and J. C. Hou, "Mixed mode simulation for IEEE 802.11-operated WLANs: integration of packet mode and fluid model based simulation," Proceedings of 2005 International Symposium on Performance Evaluation of Computer and Telecommunication Systems, July 2005.

G. He and J. Hou, "Tracking targets with quality in wireless sensor networks," Proceedings of IEEE International Conference on Network Protocols, November 2005.

S. Cheung and K. N. Levitt, "A formal-specification based approach for protecting the Domain Name System," Proceedings of the IEEE International Conference on Dependable Systems and Networks, New York, NY, June 2000, pp. 641-651.

M. Tylufki and K. Levitt, "Mitigating distributed denial of service attacks using a proportional-integral-derivative controller," Proceedings of Recent Advances in Intrusion Detection, Pittsburg, PA, September 2003.

I. Balepin, S. Maltsev, J. Rowe and K. Levitt, "A specification-based intrusion detection for automated response," Proceedings of Recent Advances in Intrusion Detection, Pittsburg, PA, September 2003.

T. Song, C. Ko, C. H. Tseng, P. Balasubramanyam, A. Chaudhary, K. N. Levitt, "Formal reasoning about a specification-based intrusion detection for dynamic auto-configuration protocols in ad hoc networks," Proceedings of Formal Aspects of Security and Trust, 2005.

C. H. Tseng, T. Song, P. Balasubramanyam, C. Ko and K. Levitt, "A specification-based intrusion detection model for OLSR," Proceedings of Recent Advances in Intrusion Detection, 2005.

B. G. Chun and M. Baker, "Evaluation of packet scheduling algorithms in mobile ad hoc networks." ACM Mobile Computing and Communications Review, vol. 6, issue 3, July 2002, pp. 36-49.

S. Marti, T. Giuli, K. Lai and M. Baker, "Mitigating routing misbehavior in mobile ad hoc networks," Proceedings of Mobicom 2000, Boston, MA, August 2000.

C. Workshop Papers

P. Narasimhan, L. E. Moser and P. M. Melliar-Smith, "Strong replica consistency for fault-tolerant CORBA applications," Proceedings of the IEEE 6th Workshop on Object-Oriented Real-Time Dependable Systems, Rome, Italy, (January 2001), 10-17.

S. Oberg, L. Tewksbury, L. E. Moser and P. M. Melliar-Smith, "Online upgrades for CORBA and EJB/J2EE," Proceedings of the Workshop on Dependable Middleware-Based Systems Washington, DC (June 2002), G38-G48.

L. E. Moser, P. M. Melliar-Smith and L. A. Tewksbury, "Online upgrades become standard," Proceedings of the IEEE Workshop on Dependable Online Upgrading of Distributed Systems, Oxford, England (August 2002), 982-988.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Deterministic scheduling for multithreaded replicas," Proceedings of the IEEE 10th International Workshop on Object-Oriented Real-Time Dependable Systems, Sedona, AZ (February 2005), 74-81.

M. Panahi, T. Harmon and R. Klefstad, "Adaptive techniques for minimizing middleware memory footprint for distributed, real-time, embedded systems," Proceedings of the IEEE Computer Communications Workshop, 2003, pp. 54-58.

R. Klefstad, M. Despande, C. O'Ryan and D. C. Schmidt, "ZEN: Implementing real-time CORBA with real-time Java," Proceedings of the Real-Time and Embedded Distributed Object Systems Workshop, Object Management Group, 2002.

J. Suzuki and T. Suda, "Middleware support for super distributed autonomic services in pervasive networks," Proceedings of the Workshop on Service Oriented Computing in the IEEE SAINT Symposium, 2004.

T. Nakano and T. Suda, "Evolution and adaptation of distributed network services," Proceedings of the Workshop on Service Oriented Computing in the IEEE International Symposium on Applications and the Internet, January 2004.

M. Moore and T. Suda, "Adaptable peer-to-peer discovery of objects that match multiple keywords," Proceedings of the Workshop on Service Oriented Computing in the IEEE International Symposium on Applications and the Internet, January 2004.

K. Fujii and T. Suda, "Component service models with semantics (CoSMoS): A new component model for dynamic service composition," Proceedings of the Workshop on Service Oriented Computing in the IEEE International Symposium on Applications and the Internet, January 2004.

R. Egashira, A. Enomoto and T. Suda, "Distributed and adaptive discovery using preference," Proceedings of the Workshop on Service Oriented Computing in the IEEE International Symposium on Applications and the Internet, January 2004.

A. Yamamoto, D. Asahara, T. Itao, S. Tanaka and T. Suda, "Distributed page rank: A distributed reputation model for open peer-to-peer networks," Proceedings of the Workshop on Service Oriented Computing in the IEEE SAINT Symposium, 2004.

H. Sasaki, H. Iwasaki and T. Suda, "Simulation of information propagation for vehicles in physical communication network models," Proceedings of the Workshop on Service Oriented Computing in the IEEE SAINT Symposium, 2004.

J. Suzuki and T. Suda, "The bio-networking platform: An autonomic agent platform for pervasive networks," Proceedings of the 2nd IPSJ Workshop on Ubiquitous Computing, Kyoto, Japan, November 2003.

R. Egashira and T. Suda, "Distributed discovery using preference for query hit," Proceedings of the IEEE Computer Communications Workshop, October 2003.

J. Lu and T. Suda, "Coverage-aware self-scheduling in sensor networks," Proceedings of the IEEE Computer Communication Workshop, October 2003.

M. Yang, Y. Huang, J. Kim, M. Lee, T. Suda and M. Daisuke, "An end-to-end QoS framework with on-demand bandwidth reconfiguration," Proceedings of the IEEE Computer Communication Workshop, October 2003.

Y. Pan, M. Lee, J. B. Kim and T. Suda, "Smooth handoff scheme for stream media with bandwidth disparity in wireless cells," Proceedings of the IEEE Computer Communications Workshop, October 2003.

Y. Pan, M. Lee, J.B. Kim and T. Suda, "An end-to-end multi-path smooth handoff scheme for stream media," Proceedings of the ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots, September 2003.

T. Itao, S. Tanaka and T. Suda, "Implementing self-organized applications that adapt to user behaviors," Proceedings of the IPSJ Entertainment Computing Workshop, January 2003.

J. Suzuki, T. Nakano, K. Fujii, N. Ikeda and T. Suda, "Dynamic reconfiguration of network applications and middleware systems in the bio-networking architecture," Proceedings of the IEEE Workshop on Large Scale Real-Time and Embedded Systems, December 2002.

J. Lu, W. Zeng, M. Yang and T. Suda, "Sensor network research in disaster scenarios," Proceedings of the IEEE Workshop on Large Scale Real-Time and Embedded Systems, December 2002.

J. Suzuki and T. Suda, "Middleware support for disaster response infrastructure," Proceedings of the First IEEE Workshop on Disaster Recovery Networks, June 2002.

J. Suzuki and T. Suda, "Adaptive behavior selection of autonomous objects in the bio-networking architecture," Proceedings of the 1st IEEE Annual Symposium on Autonomous Intelligent Networks and Systems, May 2002.

K. Fujii and T. Suda, "Loose interface definition: An extended interface definition for dynamic service composition," Proceedings of the First IEEE Annual Symposium on Autonomous Intelligent Networks and Systems, May 2002.

M. Moore and T. Suda, "A decentralized and self-organizing discovery mechanism," Proceedings of the 1st IEEE Annual Symposium on Autonomous Intelligent Networks and Systems, May 2002.

N. Yamashita and T. Suda, "An equilibrium model of a self-organizing network architecture," Proceedings of the IEICEJ Next Generation Networks Workshop, December 2001.

T. Nakamura, M. Matsuo, T. Itao and T. Suda, "A study about the system control of an adaptive networking architecture for service emergence," Proceedings of the IEICEJ Next Generation Networks Workshop, December 2001.

S. Sugizaki and T. Suda, "Evolution simulations for the bio-networking architecture," Proceedings of the IEICEJ Next Generation Networks Workshop, December 2001

T. Suda, "Biologically inspired network service architecture: BioNet," Proceedings of the Santa Fe Institute Workshop on The Internet as a Large-Scale Complex System, March 2001.

J. Suzuki and T. Suda, "Bionet project overview: Applying biological concepts and mechanisms for designing adaptable, scalable and survivable communication software," Proceedings of the JSPS Workshop on Applied Information Technology for Science, 2001.

S. Song and T. Suda, "Security consideration in the bio-networking architecture," Proceedings of the JSPS Workshop on Applied Information Technology for Science, 2001.

T. Itao, T. Suda, T. Nakamura and M. Matsuo, "Adaptive networking architecture for service emergence," Proceedings of the Workshop on Highly Distributed Systems in the IEEE SAINT Symposium, 2001.

T. Itao, T. Nakamura, M. Imada, M. Matsuo, T. Suda and T. Aoyama, "The model and design of cooperative interaction for service composition," Proceedings of IPSJ DCOMO Workshop, June 2001, pp. 7-12.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Multicasting in ad hoc networks in the context of multiple channels and multiple interfaces," Proceedings of the International Workshop on Localized Communication and Topology Protocols for Ad hoc Networks, November 2005, Washington, D.C.

M. A. Spohn and J. J. Garcia-Luna-Aceves, "Improving broadcast operations in ad hoc networks using two-hop connected dominating sets," Proceedings of IEEE Globecom 2004 Workshop on Wireless Ad Hoc and Sensor Networks, Dallas, TX, November/December 2004.

S. Roy and J. J. Garcia-Luna-Aceves, "Node-centric hybrid routing for ad hoc networks," Proceedings of the 10th IEEE/ACM Workshop on Mobility and Wireless Access, October 2002.

Y. Wang and J. J. Garcia-Luna-Aceves, "Collision avoidance in multi-hop ad hoc networks," Proceedings of the 10th IEEE/ACM Workshop on Mobility and Wireless Access, Fort Worth, TX, October 2002.

J. J. Garcia-Luna-Aceves, "Flow-oriented protocols for scalable wireless networks," Keynote Talk, The Fifth ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems, Atlanta, GA, September 2002.

Y. Wang and J. J. Garcia-Luna-Aceves, "A new hybrid channel access scheme for ad hoc networks," Proceedings of the IFIP First Annual Mediterranean Ad Hoc Networking Workshop, Sardegna, Italy, September 2002.

H. Zhang and J. Hou, "A scheduling algorithm for variable rate coded voice in Bluetooth networks," Proceedings of the Fifth ACM International Workshop on Wireless Mobile Multimedia, September 2002.

H. Zhang and J. C. Hou, "Maintaining sensing coverage and connectivity in large sensor networks," Proceedings of the NSF Workshop on Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless, and Peer-to-Peer Networks, Fort Lauderdale, FL, February 2004.

T. Song, J. Alves-Foss, C. Ko, C. Zhang and K. Levitt, "Using ACL2 to verify security properties of specification-based intrusion detection systems," Proceedings of the Fourth International Workshop on the ACL2 Theorem Prover and Its Applications, 2003.

C. Y. Tseng, P. Balasubramanyam, C. Ko, R. Limprasittiporn, J. Rowe and K. Levitt , "A specification-based intrusion detection system for AODV," Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks, 2003.

P. Porras, L. Briesemeister, K. Skinner, K. Levitt, J. Rowe, and Y. C. A. Ting, "A hybrid quarantine defense," Proceedings of WORM, CCS Workshop on Rapid Malcode, October 2004, Washington, D.C.

D. Sterne, P. Balasubramanyam, D. Carman, B. Wilson, R. Talpade, C. Ko, R. Balupari, C. H. Tseng, T. Bowen, K. Levitt and J. Rowe, "A general cooperative intrusion detection architecture for MANETs," Proceedings of the 3rd IEEE International Information Assurance Workshop, 2005.

S. H. Wang, C. Y. Tseng, C. Ko, M. Gertz and K. Levitt, "A general automatic response model for MANETs," Proceedings of the First IEEE International Workshop on Next Generation Wireless Networks 2005.

D. Encyclopedia Articles and Book Chapters

V. Kalogeraki, P. M. Melliar-Smith and L. E. Moser, "Managing real-time distributed multimedia applications," *Multimedia Networking: Technology, Management and Applications*, ed. Mahbubur Rahman Syed, Idea Group Publishing, Hershey, PA (January 2002), 1-16.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "Fault tolerance for distributed and networked applications," *Encyclopedia of Information Science and Technology* (January 2005).

P. M. Melliar-Smith and L. E. Moser, "Mobile multimedia computing for commerce," *Encyclopedia of Multimedia Technology and Networking*, Idea Group Publishing (June 2005).

S. Munjee, N. Surendran, Y. Krishnamurthy and D. C. Schmidt, "The design and performance of a CORBA audio/video streaming service," *Multimedia Networking: Technology, Management, and Applications*, edited by Mahbubur Syed, Idea Group Publishing, Hershey, PA, 2001.

C. O'Ryan, F. Kuhns, D. C. Schmidt and J Parsons, "Applying patterns to develop a pluggable protocols framework for ORB middleware," *Design Patterns in Communications*, Linda Rising, ed., Cambridge University Press, 2000.

A. Krishna, D. C. Schmidt, R. Klefstad and A. Corsaro, "Real-time CORBA middleware, Middleware for Communications," ed. Qusay Mahmoud, John Wiley & Sons, New York, 2003.

T. Suda, T. Itao and M. Matsuo, "The bio-networking architecture: A biologically inspired approach to the design of scalable, adaptive, and survivable/adaptable network applications," in *The Internet as a Large Complex System*, Santa Fe Institute Book Series, Oxford University Press, 2005.

L. Bao and J. J. Garcia-Luna-Aceves, "Distributed channel access scheduling for ad hoc networks," *Algorithms and Protocols for Wireless and Mobile Networks* (A. Boukerche, ed.), CRC/Hall Publisher, 2004.

J. J. Garcia-Luna-Aceves and Y. Wang, "MAC protocols for ad hoc networks," Chapter 2, *Ad Hoc Networks: Technologies and Protocols*, Kluwer, 2004.

J. J. Garcia-Luna-Aceves and S. Roy, "Node-centric hybrid routing for wireless internetworking," *Mobile and Wireless Internet* (K. Makki, et al eds.), Chapter 8, Kluwer, 2003.

R. Zheng, J. C. Hou and N. Li, "Power management and control in wireless networks," *Ad Hoc and Sensor Networks*, Y. Pan and Y. Xiao (eds), Nova Science Publishers, Fall 2005.

H. Zhang and J. C. Hou, "Maintaining sensing coverage and connectivity in large sensor networks," *Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless and Peer-to-Peer Network*, Jie Wu (ed), Chapter 28, Auerbach Publications, January 2005, pp. 453-475.

J. C. Hou, I. Stojmenovic and Ning Li, "Topology construction and maintenance in sensor networks," *Handbook of Sensor Networks: Algorithms and Architectures*, I. Stojmenovic (ed), John Wiley and Sons, 2005, pp. 311-342.

N. Li and J. C. Hou, "Topology control in wireless sensor networks," *Combinatorial Optimization in Communication Networks*, D. Z. Du, M. Cheng and Y. Li (eds), Kluwer Publishers, Spring 2005.

W. P. Chen and J. C. Hou, "Data gathering and fusion in sensor networks," *Handbook of Sensor Networks: Algorithms and Architectures*, I. Stojmenovic (ed), John Wiley and Sons, 2005, pp. 493-526.

E. Standards Documents

P. M. Melliar-Smith and L. E. Moser, "Online Upgrades," Object Management Group, mars/2002-05-01 (May 2002).

L. E. Moser, P. M. Melliar-Smith, et. al., "Service Availability Forum Application Interface Specification," Service Availability Forum (April 2003).

S. Sameshima, J. Suzuki, S. Steglich and T. Suda, "Platform Independent Model (PIM) and Platform Specific Model (PSM) for Super Distributed Objects," Object Management Group, final adopted specification, OMG document number: dtc/03-09-01, Sept. 2003.

F. In Press

P. Dasgupta, L. E. Moser and P. M. Melliar-Smith, "Dynamic pricing algorithms for online sellers using intelligent agents," Encyclopedia of E-Commerce, E-Government and Mobile Applications.

L. E. Moser and P. M. Melliar-Smith, "Web Services," Encyclopedia of E-Commerce, E-Government and Mobile Applications.

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "High availability and data consistency for three-tier enterprise applications," Encyclopedia of E-Commerce, E-Government and Mobile Applications.

P. Dasgupta, L. E. Moser and P. M. Melliar-Smith, "Maximizing welfare through cooperative negotiation in a multi-agent Internet economy," International Journal of Information Technology and Decision Making.

T. Nakano and T. Suda, "Applying biological principles to the design of network services," Elsevier, Applied Soft Computing Journal.

K. Fujii and T. Suda, "Semantics-based dynamic service composition," IEEE Journal on Selected Areas in Communications (JSAC), Special Issue on Autonomic Communication Systems.

M. Yang, Y. Huang, J. Kim, M. Lee, T. Suda and M. Daisuke, "A traffic adaptive QoS framework with distributed admission control," Computer Communication Journal, Special Issue on End-to-End Quality of Service Differentiation, Elsevier.

A. Yahaya and T. Suda, "iREX: Inter-domain QoS automation using economics," IEEE Consumer Communications and Networking Conference, 2006.

Y. Wang and J. J. Garcia-Luna-Aceves, "Performance of directional collision avoidance in multihop wireless networks," Special Issue on Wireless Networks and Pervasive Computing, Journal of Pervasive Computing and Communications, Issue 1, 2006.

H. Zhang and J. C. Hou, "Is deterministic deployment worse than random deployment for wireless sensor networks," Proceedings of IEEE INFOCOM, March 2006.

H. Kim, H. Lim, and J. C. Hou, "Accelerating simulation of large-scale IP networks: an network invariant preserving approach," Proceedings of IEEE INFOCOM 2006, March 2006.

C. Hu, R. Zheng and J. C. Hou, "A microscopic study of power management in IEEE 802.11 wireless networks," International Journal of Wireless and Mobile Computing, Special Issue on Medium Access Control for WLANs, WPANs, Ad Hoc Networks, and Sensor Networks.

R. Zheng, J. C. Hou, and L. Sha, "Optimal block design for asynchronous wakeup and its applications in multi-hop wireless networks," IEEE Transactions on Mobile Computing.

N. Li and J. C. Hou, "Topology control in heterogeneous wireless networks: problems and solutions," IEEE/ACM Transactions on Networking.

N. Li and J. C. Hou, "BLMST: A decentralized, power-efficient broadcast algorithm for wireless sensor networks," ACM/Baltzer Wireless Networks.

N. Li and J. C. Hou, "FLSS: A fault-tolerant topology control algorithm for wireless networks," IEEE Transactions on Parallel and Distributed Systems.,

A. Sobeih et al., "J-Sim: A network simulation and emulation environment for wireless sensor networks," IEEE Wireless Communications Magazine.

G. He and J. C. Hou, "On sampling self-similar Internet traffic," Computer Networks Journal.

J. C. Hou and P. R. Kumar, "Network modeling and simulation," Editorial for the Special Issue of Network Modeling and Simulation, Computer Networks Journal.

C. Hu, H. Kim, J. C. Hou and S. Shankar, "Provisioning quality-controlled medium access in ultra wide band WPANs," Proceedings of IEEE INFOCOM, March 2006.

H. Lim, L. Kung, H. Luo and J. C. Hou, "Use of singular value decomposition for zero-configuration, robust indoor localization: theory and experimentation," Proc. IEEE INFOCOM 2006, March 2006).

G. In Submission

V. Kalogeraki, L. E. Moser and P. M. Melliar-Smith, "Resource management using multiple feedback loops in soft real-time distributed object systems."

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "A reservation-based extended transaction protocol for coordination of Web services."

W. Zhao, L. E. Moser and P. M. Melliar-Smith, "A reservation-based transaction protocol."

L. E. Moser, P. M. Melliar-Smith and W. Zhao, "Making Web Services dependable."

J. A. Colmenares, S. Gorappa, M. Panahi and R. Klefstad, "A component framework for real-time Java."

H. In Preparation

M. Schuricht, Z. Davis, Y. X. Hu, S. Prasad, P. M. Melliar-Smith and L. E. Moser, "Managing multiple speech-enabled applications in a mobile handset."

S. Prasad, Z. Davis, Y. X. Hu, M. Schuricht, P. M. Melliar-Smith and L. E. Moser, "Location-based services for mobile voice-enabled devices."

Z. Davis, Y. X. Hu, S. Prasad, M. Schuricht, P. M. Melliar-Smith and L. E. Moser, "A personal handheld multi-modal shopping assistant."

Y. X. Hu, Z. Davis, S. Prasad, M. Schuricht, P. M. Melliar-Smith and L. E. Moser, "Speech-enabled Web Services on mobile devices."

A. Y. Ting, C. Ko, J. Rowe and K. Levitt, "Specification-based network intrusion detection: Detecting attacks on DHCP."

P. Balasubramanyam and K. Levitt, "Topology monitoring of mobile ad hoc networks."

C. H. Tseng, K. Levitt, et al, "DEMEN: A distributed event-driven message exchange model for intrusion detection in MANET."

C. G. Senthilkumar, K. Levitt, et al, "An API for worm defense."

H. Ph.D. Dissertations and M.S. Theses

L. Tewksbury, "Live Software Upgrades Using Replication of CORBA Objects," Ph.D. Dissertation, Department of Electrical and Computer Engineering, University of California, Santa Barbara, Spring 2001.

N. Narasimhan, "Transparent Fault Tolerance for Java Remote Method Invocation," Ph.D. Dissertation, Department of Electrical and Computer Engineering, University of California, Santa Barbara, Spring 2001.

P. Dasgupta, "Mobile Software Agent Enabled E-Commerce: System Design and Profit Maximizing Algorithms," Ph.D. Dissertation, Department of Electrical and Computer Engineering, University of California, Santa Barbara, Spring 2001.

W. Zhao, "Unification of Replication with Transaction Processing," Ph.D. Dissertation, Department of Electrical and Computer Engineering, University of California, Santa Barbara, Spring 2002.

S. Gwalani, "Intrusion Detection in Wireless Ad hoc Networks," Department of Computer Science, University of California, Santa Barbara, Summer 2003.

K. Srinivasan, "AODVSTAT - An Intrusion Detection System for Ad hoc Networks," Department of Computer Science, University of California, Santa Barbara, Fall 2004.

M. Panahi, "Customizing Middleware for Footprint Reduction in Memory-Constrained Systems", M.S. Thesis, School of Information and Computer Science, University of California, Irvine.

M. A. Spohn, "Domination in Graphs in the Context of Mobile Ad Hoc Networks", Ph.D. Dissertation, Computer Science, University of California, Santa Cruz, July 2005.

M. Mosko, "Routing in Mobile Ad Hoc Networks", Ph.D. Dissertation, Computer Engineering, University of California, Santa Cruz, June 2004.

Y. Wang, "Medium Access Control in Ad Hoc Networks with Omni-directional and Directional Antennas", Ph.D. Dissertation, Computer Engineering, University of California, Santa Cruz, CA, June 2004.

S. Roy, "On-Demand Link-State Routing in Ad-Hoc Networks", Ph.D. Dissertation, Computer Engineering, University of California, Santa Cruz, CA, June 2003.

L. Bao, "Neighbor-Aware Control in Ad Hoc Networks", Ph.D. Dissertation, Computer Science, University of California, Santa Cruz, CA, December 2002.

M. Tylutki, "Network-based Response Framework and Implementation," Ph.D. Dissertation, Department of Computer Science, University of California, Davis, August 2005.

H. Y. Tyan, "Design, realization and evaluation of a component-based compositional software architecture for network simulation," Ph.D. Dissertation, Department of Electrical Engineering, Ohio State University, December 2001.

Y. Gao, "Congestion control for the next generation Internets," Ph.D. Dissertation, Department of Electrical Engineering, Ohio State University, September 2002.

G. He, "Exploitation of long-range dependency in Internet traffic in resource and traffic control," Ph.D. Dissertation, Department of Computer Science, University of Illinois at Urbana Champaign, August 2004.

N. Li, "Traffic control in multi-hop ad-hoc networks," Ph.D. Dissertation, Department of Computer Science, University of Illinois at Urbana Champaign, August 2005.

H. Zhang, "On understanding fundamental performance limits of wireless sensor networks," Ph.D. Dissertation, Department of Computer Science, University of Illinois at Urbana Champaign, August 2005.

8. Presentations

Michael Melliar-Smith, "Protocol Engineering Research Center" and "ARCH: Clustered Routing for Ad-hoc Mobile Networks," MURI Kickoff Meeting, Berkeley, CA, May 2000.

Michael Melliar-Smith, "The Future of Middleware," WORDS Workshop, San Diego, CA, January 2002.

Michael Melliar-Smith, "Protocol Engineering Research Center," MURI PI Meeting, Santa Barbara, CA, August 2002.

Michael Melliar-Smith, "Protocol Engineering Research Center" and "Fault Tolerance for CORBA Using the Pluggable Protocols Framework," MURI PI Meeting, Irvine, CA, October 2002.

Michael Melliar-Smith, "Protocol Engineering Research Center," MURI Progress Review Meeting, AFOSR Arlington, VA, December 2002.

Michael Melliar-Smith, "Application Interface Specification," Service Availability Forum, Batvik, Finland, May 2003.

Michael Melliar-Smith, "Real-Time Fault-Tolerant CORBA," Object Management Group, Arlington, VA, June 2003.

Michael Melliar-Smith, "Real-Time Fault Tolerance," IEEE 37th Hawaii International Conference on System Sciences, Kona, HI, January 2004.

Michael Melliar-Smith, "The Heisenbug Problem," IEEE International Conference on Dependable Systems and Networks, Florence, Italy, June 2004.

Michael Melliar-Smith, "Progress in Real-Time Fault Tolerance," IEEE Symposium on Reliable Distributed Systems, Florianopolis, Brazil, October 2004.

Louise Moser, "FTMP: A Fault Tolerant Multicast Protocol," MURI Kickoff Meeting, Berkeley, CA, May 2000.

Louise Moser, "Fault Tolerant CORBA Tutorial," Fourth European Research Seminar on Advances in Distributed Systems," Bertinoro, Italy, May 2001.

Louise Moser, "State Synchronization and Recovery for Strongly Consistent Replicated CORBA Objects," IEEE 2001 International Conference on Dependable Systems and Networks, Goteberg, Sweden, June 2001.

Louise Moser, "Fault Tolerant CORBA Tutorial," IEEE Conference on Dependable Systems and Networks, Goteburg, Sweden, June 2001.

Louise Moser, "Increasing the Reliability of Three-Tier Applications," IEEE 12th International Symposium on Reliability Engineering, Hong Kong, November 2001.

Louise Moser, "Multicast Protocols for Fault-Tolerant Systems," MURI PI Meeting, Santa Barbara, CA, August 2002.

Louise Moser, "Fault Tolerance for CORBA, Java and Operating Systems," Defense Interoperability Initiative, Common Operating Environment, Real-Time Advisory Group, CMU Software Engineering Institute, Pittsburgh, PA, August 2002.

Louise Moser, "Online Upgrades Become Standard," IEEE COMPSAC Workshop on Dependable On-line Upgrading of Distributed Systems, Oxford, England, August 2002.

Louise Moser, "Secure Multicast Protocols for CORBA" and "Testing and Analysis of Distributed Systems and Protocols," MURI Workshop, Irvine, CA, October 2002.

Louise Moser, "Application Interface Specification," Service Availability Forum Meeting, Helsinki, Finland, December 2002.

Louise Moser, "Checkpointing and Logging for Intrusion Analysis and Recovery," IEEE International Conference on Dependable Systems and Networks, Florence, Italy, June 2004.

Michael Melliar-Smith and Louise Moser, "Fault Tolerant CORBA Tutorial," Object Management Group Meeting, Burlingame, CA, September 2000.

Michael Melliar-Smith and Louise Moser, "Fault Tolerant CORBA Tutorial," Lucent CORBA Forum, Naperville, IL, November 2000.

Michael Melliar-Smith and Louise Moser, "Online Upgrades for CORBA," Object Management Group, Anaheim, CA, January 2002.

Michael Melliar-Smith and Louise Moser, "Service Availability Architecture and Specification Proposal," Service Availability Forum, Atlanta, GA, April 2002.

Michael Melliar-Smith and Louise Moser, "Fault Tolerant CORBA and Carrier-Grade Linux," Lucent Technologies, Naperville, IL, September 2002.

Michael Melliar-Smith and Louise Moser, "Reliable Ordered Multicast Inter-ORB Protocol," Object Management Group, Anaheim, CA, November 2002.

Raymond Klefstad, "Reconfigurable and Adaptive Middleware for Distributed, Real-time, Embedded Systems," Center for Embedded Computer Systems Symposium on RASC technology, September 30, 2003.

Tatsuya Suda has given invited presentations at numerous companies, including NTT, NTT Docomo, Denso IT Laboratory, Hitachi, Hitachi America, KDDI, and NICT.

Garcia-Luna-Aceves was a Plenary Speaker at the 4th International Conference on Ad Hoc Networks and Wireless, Cancun, Mexico, October 2005.

9. Other Professional Activities

Michael Melliar-Smith and Louise Moser have participated in the standardization activities of the Object Management Group, including the specifications for Fault Tolerant CORBA, Portable Interceptors, Unreliable Multicast, Dynamic Scheduling, and Online Upgrades.

Michael Melliar-Smith and Louise Moser have participated in the standardization activities of the Service Availability Forum. Louise Moser has served as editor of the Application Interface Specification for the Service Availability Forum.

Michael Melliar-Smith and Louise Moser have participated in the standardization activities of the Java Community Process, JSR 117 Continuous Availability.

Michael Melliar-Smith has served as an invited panel member of the Panel on Real-Time Fault Tolerance at the IEEE 37th Hawaii International Conference on System Sciences, Kona, HI (January 2004). He has also served as an invited panel member of the Panel on Progress in Real-Time Fault Tolerance at the IEEE Symposium on Reliable Distributed Systems, Florianopolis, Brazil (October 2004). He has also served on the Publicity and Advisory Committee of the IEEE Workshop on Object-Oriented Real-Time Dependable Systems (2005).

Louise Moser has served on the Conference Program Committees of the IFIP Middleware Conference (2000, 2001), IEEE Symposium on Reliable Distributed Systems (2000, 2002, 2003), Technology of Object-Oriented Languages and Systems (TOOLS) (2001), IEEE Conference on Dependable Systems and Networks, (2002, 2003), IEEE Workshop on Object-Oriented Real-Time Dependable Systems (2002), IEEE International Symposium on Object-Oriented Real-Time Distributed Computing, (2000, 2001, 2002). She has also served on the Publicity and Advisory Committee of the IEEE Workshop on Object-Oriented Real-Time Dependable Systems (2005).

Louise Moser has served as a Session Chair for the Fourth European Research Seminar on Advances in Distributed Systems Bertinoro, Italy (May 2001) and a Session Chair for the IEEE International Conference on Services Computing (July 2005).

Louise Moser has served as Associate Editor of the IEEE Transactions on Computers (2000, 2001). She has also served as a member of the Scientific Advisory Board, Dependability Foundations for Information Infrastructures - Network of Excellence (2003, 2004).

Elizabeth Belding-Royer served as Technical Program Co-Chair of ACM MobiCom in 2005, and as Technical Program Co-Chair of IEEE SECON in 2005.

Elizabeth Belding-Royer served as an Editorial Board Member of the Ad Hoc Network Journal, 2002-2005, and as an Editorial Board Member of the Transactions on Mobile Computing, 2005-present.

Mary Baker was elected by the membership to be vice chair of the ACM Special Interest Group on Mobility (SIGMOBILE). In addition, she served as co-chair of the technical program committee of the new annual ACM/USENIX conference on higher-level issues in mobile systems. The first MobiSys conference took place in May 2003 and is a premier forum for publishing mobile systems work. She has also served on many program committees and editorial boards in the area of mobility and is also one of the founding editorial board members of the new IEEE Pervasive Computing magazine.

Raymond Klefstad presented RTZen at Java day, Raytheon, Bedford, MA, July 2005.

Tatsuya Suda held a workshop on Service Oriented Computing at the IEEE Symposium on Applications and the Internet Conference, 2004.

Garcia-Luna-Aceves has served as General Chair of the IEEE SECON 2005 Conference, September 2005.

Garcia-Luna-Aceves has served as Program Co-Chair of the ACM MobiHoc 2002 Conference, June 2002.

Jennifer Hou has served (or will serve) as the co-chair of the technical program committee of IEEE Real-time Technology and Application Symposium in June 2000, 3rd IEEE Information Processing in Sensor Networks in April 2004, 1st Wireless Internet Conference in July 2005, 4th IEEE International Conference on Mobile Ad-hoc and Sensor Systems in October 2006, and 29th IEEE INFOCOM in March 2008. She has also served as the vice chair of the technical program committee of the IEEE International Conference on Distributed Computing Systems in June 2002, IEEE International Conference on Parallel and Distributed Systems for the wireless and mobile computing track in June 2004, and IEEE 25th Real-Time Systems Symposium for the real-time communications and sensor network track in December 2004. In addition, she has served on the technical program committees for major networking and real-time conferences such as IEEE INFOCOM, IEEE ICNP, IEEE ICDCS, ACM MobiCom, ACM MobiHoc, ACM/IEEE IPSN, IEEE RTSS, and IEEE RTAS.

Jennifer Hou has served on the editorial boards of IEEE Transactions on Wireless Communications, IEEE Transactions on Parallel and Distributed Systems, ACM Transactions on Sensor Networks, IEEE Wireless Communication Magazine, Computer Networks, ACM Wireless Networks, and Foundations and Trends in Networking.

Jennifer Hou has guest edited a special issue on Wireless Sensor Networks: Theory and Systems in IEEE Wireless Communications Magazine and a special issue on Network Modeling and Simulation in the Computer Networks Journal. She has also been invited by Springer to edit a book titled Wireless Mesh Networks: Design, Evaluation, and Deployment.

Jennifer Hou has served as a panelist on the panel Replacing Copper by Radio? What are Research and Business Opportunities? at IEEE INFOCOM 2004 in Hong Kong, China, March 2004 and has also served as a panelist on the panel Challenges in Mobile/Ubiquitous Networking/Services at ACM MobiQuitous 2005 in San Diego, July 2005. In additional, she has given a keynote presentation at the 2005 International Symposium on Communications, Kaohsiung, Taiwan.

10. Technology Transfer

The principal investigators of this AFOSR MURI project have relied on several different methods for technology transfer into industry. These methods include meetings with collaborators in commercial organizations, release of public domain software, standardization activities, and formation of startup companies.

Michael Melliar-Smith, Louise Moser, Priya Narasimhan, Lauren Tewksbury, Wenbing Zhao, Ruppert Koch, and Nitya Narasimhan have all participated in a startup company, Eternal Systems, which is developing software products for high availability and fault tolerance, based in part on the research of this MURI project and of prior DARPA projects at UCSB. Specific technology transferred to Eternal Systems

has included Fault Tolerant CORBA technology, consistent checkpointing mechanisms, TCP connection failover technology, and low-latency real-time multicast and group membership protocols. Patents related to these technologies have been assigned to Eternal Systems. Louise Moser, Michael Melliar-Smith and their students have participated in the development and release of the FTORB Fault Tolerant CORBA product and the Duration checkpointing/restart product for Eternal Systems.

Louise Moser and Michael Melliar-Smith participated in the development of standards for the Object Management Group, including Fault Tolerant CORBA, Online Software Upgrades, and Dynamic Scheduling. CORBA is used in military and defense systems.

Louise Moser, Michael Melliar-Smith participated in the development of the Service Availability Forum Application Interface Specification. The Service Availability Forum is a consortium of about 30 companies, including Intel, Sun, IBM, HP, Nokia, Motorola, and Fujitsu-Siemens that is developing standards for high availability and fault tolerance middleware for telecommunications applications. The Application Interface Specification defines interfaces for availability management, cluster membership service, checkpoint service, event service, message service, and lock service. Louise Moser was the editor of the Application Interface Specification.

Douglas Schmidt has released TAO as freely available open source software at <http://www.ece.uci.edu/~schmidt/TAO.html>

Tatsuya Suda has been working actively in the Object Management Group (OMG) to reflect the Bio-Networking platform in the OMG Super Distributed Objects specification. A detailed list of artifacts and software is included below.

Jennifer Hou has released JavaSim as freely available open source at <http://www.javasim.org>.

11. Patent Disclosures

During the period of this grant, Michael Melliar-Smith and Louise Moser has had several patents granted, and filed several other patent applications, related to Online Upgrades, High Availability and Fault Tolerance, in particular:

- Live Upgrade Process for Object-Oriented Programs; 6,360,363, granted March 19, 2002.
- Fault Tolerance for Computer Programs that Operate over a Communication Network; 20020099973, filed October 26, 2001; 6,922,792 granted July 26, 2005.
- Consistent Message Ordering for Semi-Active and Passive Replication; 20040103342, filed July 29, 2003; 6,928,577 granted August 9, 2005
- Transparent Consistent Active Replication of Multithreaded Application Programs; 20040078617, filed March 24, 2003.
- Transparent Consistent Semi-Active and Passive Replication of Multithreaded Application Programs; 20040078618, filed March 25, 2003.
- Transparent TCP Connection Failover; 20040268175, filed June 11, 2003.
- Consistent Asynchronous Checkpointing of Multithreaded Application Programs Based on Semi-active or Passive Replication; 20050034014, filed August 30, 2003.
- Consistent Asynchronous Checkpointing of Multithreaded Application Programs Based on Active Replication, filed August 30, 2003.
- Consistent Time Service for Fault-Tolerant Distributed Systems; 20040205372, filed January 2, 2004.

These patent applications and patents have been transferred and assigned to Eternal Systems.

Tatsuya Suda has had granted one patent and has filed another patent application, namely:

- A Framework for Creating Adaptive Network Services; 2002-175405, filed in December 2000, P3708432, granted in October 2005.
- A Framework for Information and Service Distribution; 2005-103154, filed in March 2005.

12. Software Releases

Using technology developed with support in part from this MURI project, Michael Melliar-Smith and Louise Moser helped Eternal Systems deliver its high availability and fault tolerance software products to multiple customers, including NSWC, Dahlgren:

- FTORB Fault Tolerant CORBA product
- Duration High Availability Checkpoint/Restore product

Elizabeth Belding-Royer and her students released the following software package:

- ARAN, Version 0.1, May 1, 2004

Raymond Klefstad and Douglas Schmidt and their students released the following software packages:

- ZEN, Version 01.00.01, September 19, 2003
- RTZen, Snapshot 588 , August 5, 2004
- ZEN-kit, Version 0.1, May 1, 2005

Tatsuya Suda and his students have released the following artifacts and software:

- Software for the Bio-Networking Architecture
 - Dynamic Service Composition software, <http://netresearch.ics.uci.edu/kfujii/dsc/>
 - Adaptation and evolution simulator, http://netresearch.ics.uci.edu/bionet/resources/evolution_simulator/index.html
 - Bio-Networking platform software, <http://netresearch.ics.uci.edu/bionet/resources/platform/>
- Software for Distributed Discovery in Peer-to-Peer Systems
 - Similarity-based Peer-to-Peer discovery protocol simulator, <http://netresearch.ics.uci.edu/mikemo/>
 - Preference-based Peer-to-Peer discovery protocol simulator, http://netresearch.ics.uci.edu/bionet/resources/discovery/ryota/discovery_simulator.zip
- Software for Protocols for Sensor Networks
 - Protocol for Coverage-Aware Sensor Engagement in Sensor Networks, <http://bolero.ics.uci.edu/lujun/simulators/CASE-ns2.tar.gz>
 - Protocol for Monitoring a Target from Multiple Viewing Angles in Sensor Networks, <http://bolero.ics.uci.edu/jidongw/simulators/>
 - Protocol for Protecting Targets in Sensor Networks, <http://www.ics.uci.edu/~meiy/mobilesensor.html>
 - Protocol for Data Dissemination in Sensor Networks, <http://bolero.ics.uci.edu/lujun/simulators/LWDD-ns2.tar.gz>
- Software for Protocols for QoS
 - Router Queue Management for Mixed Traffic, <http://bolero.ics.uci.edu/ypan/simulators/RQM.tar.gz>
 - Inter-domain QoS Automation using Economics, <http://bolero.ics.uci.edu/ariffin/simulators/iREX.tar.gz>

13. Honors

The paper, "End-to-end latency analysis and evaluation of a fault-tolerant CORBA infrastructure," by W. Zhao, L. E. Moser and P.M. Melliar-Smith, received the best paper award at the International Symposium on Performance Evaluation of Computer and Telecommunication Systems, 2002.

The paper, "Unreliable Byzantine fault detectors for solving consensus" by K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, received the Wilkes Award from the British Computer Society's Computer Journal, 2003.

The paper, "Design and performance of a dynamically configurable, messaging protocols framework for real-time CORBA," by R. Klefstad, S. Rao and D. C. Schmidt, received the best paper award in the Distributed Object and Component-based Software Systems part of the Software Technology Track at the 36th Annual Hawaii International Conference on System Sciences, 2003.

The paper, "Service emergence based on relationship among self-organizing entities," by T. Itao, T. Nakamura, M. Matsuo, T. Suda and T. Aoyama won the best paper award at the IEEE/IPSJ Symposium on Applications and the Internet, 2002.

The paper, "A method of service composition using service attributes," by M. Imada, Y. Katayama, M. Matsuo and T. Suda received the best paper award at the IEICEJ Assurance Symposium, 2001.

The paper, "Coverage aware sensor engagement in dense sensor networks," by J. Lu, L. Bao and T. Suda, won the best paper award at the IFIP International Conference on Embedded and Ubiquitous Computing, 2005.

Louise Moser received the Service Availability Forum Service Award, for her standards work as editor of the Application Interface Specification in 2003.

Elizabeth Belding-Royer was recognized as one of the Top 100 World Young Investigators by MIT Technology Review in 2002.

Elizabeth Belding-Royer received an NSF Career Award on Mobile Network Support for Collaborative Environments in 2004.

J. J. Garcia-Luna-Aceves was elected IEEE Fellow in 2005 "for contributions to theory and design of communication protocols for network routing and channel access."

J. J. Garcia-Luna-Aceves received the Jack Baskin Chair of Computer Engineering, July 1, 2003 - June 30, 2008.

Jennifer Hou received the Lumley Research Award at Ohio State University in 2001.

Jennifer Hou received an ACM Service Award, in appreciation for her contributions to ACM as program chair for the Third International Symposium on Information Processing in Sensor Networks in 2004.